On tangles, connectivity and graph-decompositions

Dissertation

zur Erlangung der Würde des
Doktors der Naturwissenschaften

der Fakultät für Mathematik, Informatik
und Naturwissenschaften

der Universität Hamburg

vorgelegt
im Fachbereich Mathematik
von

Raphael W. Jacobs

Hamburg, 2026

Betreuer:
Reinhard Diestel (Universität Hamburg)

Gutachter:
Reinhard Diestel (Universität Hamburg)
Agelos Georgakopoulos (University of Warwick)
Paul Wollan (University of Rome 'La Sapienza')

Prüfungskommission:
Ingenuin Gasser (Universität Hamburg)
Nathan Bowler (Universität Hamburg)
Ralf Holtkamp (Universität Hamburg)
Reinhard Diestel (Universität Hamburg)
Matthias Hamann (Universität Hamburg)

Datum der Disputation: 06.07.2026

Contents

Introduction **1**

I Graph-decompositions: using tangles to capture local structure **13**

1. Canonical graph-decompositions via coverings **15**

1.1 Introduction . . . 15
1.2 Terminology . . . 22
1.3 Graph-decompositions . . . 24
1.4 Local coverings . . . 34
1.5 Displaying global structure . . . 44
1.6 Local covers are accessible . . . 50
1.7 Tools and lemmas for the proof of Theorem 3 . . . 50
1.8 Proof of the main result of this chapter . . . 58
1.A Tree-decompositions in quasi-transitive graphs . . . 60

2. Canonical graph-decompositions via local separations **63**

2.1 Introduction . . . 63
2.2 Review of graph-decompositions and coverings . . . 70
2.3 Bottlenecks . . . 76
2.4 Local separations . . . 83
2.5 Displacement: a guarantee for the existence of lifts . . . 91
2.6 Links and crossing . . . 98
2.7 Local bottlenecks . . . 103
2.8 Constructing graph-decompositions from nested sets of local separations . 110
2.9 Proofs of Theorem 6 and Theorem 7 . . . 121
2.A Corners of local separations . . . 122

3. Canonical tree-decompositions of chordal graphs **125**

3.1 Introduction . . . 125
3.2 Background . . . 128
3.3 Canonical tree-decompositions into cliques . . . 131
3.4 Canonical graph-decompositions via coverings for locally chordal graphs . 138
3.5 Counterexamples to possible strengthenings of Theorem 8 . . . 142

II Tangles and concrete clusters **145**

4. On vertex sets inducing tangles **147**

4.1 Introduction 147
4.2 Preliminaries 152
4.3 Definition of 'survive' and special cases of Theorem 14 155
4.4 Proof of Theorem 15 if G has a higher-order tangle 158
4.5 Rainbow-cloud-decompositions in the absence of high-order tangles 163
4.6 RC-decompositions and separations 170
4.7 Proof of Theorem 15 if G has no higher-order tangle 174
4.8 The inductive proof method and its applications 184

5. Point sets and functions inducing tangles of set separations **189**

5.1 Introduction 189
5.2 Preliminaries 193
5.3 Resilience and locally induced orientations 196
5.4 Orientations and duality of set separations 201
5.5 Extendable tangles are induced by point sets 203

6. Efficiently distinguishing all tangles in locally finite graphs **207**

6.1 Introduction 207
6.2 Preliminaries 209
6.3 Limit separations 210
6.4 The interlacing theorem 216
6.5 Proofs of Theorems 16 and 17 222
6.6 Infinitely many rays beyond the limit 226

III Connectivity and tree-structures **229**

7. Linked tree-decompositions into finite parts **231**

7.1 Introduction 231
7.2 Preliminaries 238
7.3 Displaying ends 243
7.4 A high-level proof of the main result of this chapter 245
7.5 Tree-decomposition along critical vertex sets 250
7.6 Lifting paths and rays from torso 254
7.7 Linked tree-decompositions into rayless parts 257
7.8 Lean tree-decompositions 271
7.9 Tree-decompositions distinguishing infinite tangles 277
7.A A closer look at Theorem 23' 278

8. Counterexamples regarding linked and lean tree-decompositions 281
8.1 Introduction . . . 281
8.2 Preliminaries . . . 285
8.3 Examples 25 and 26 – negative answers to questions (i), (ii) and (iv) . . . 286
8.4 Example 27 – negative answer to question (vi) . . . 292
8.5 Upwards disjointness of adhesion sets . . . 293

9. A grid theorem for strong immersions of walls 297
9.1 Introduction . . . 297
9.2 Almost α-thin graphs and 3-centres . . . 300
9.3 Strongly immersed walls are obstructions . . . 302
9.4 Excluding strongly immersed walls . . . 306

10. The Lovász-Cherkassky theorem for locally finite graphs with ends 315
10.1 Introduction . . . 315
10.2 Tools and terminology . . . 317
10.3 Proof of Theorem 30 . . . 318

11. Menger's theorem in bidirected graphs 321
11.1 Introduction . . . 321
11.2 Preliminaries . . . 323
11.3 Counterexample to a general Menger-type theorem . . . 327
11.4 Appendages of paths . . . 329
11.5 Menger's theorem for edge-disjoint paths . . . 332
11.6 Menger's theorem for vertex-disjoint paths . . . 336
11.7 Polynomial time algorithm . . . 342

12. Hitting cycles through prescribed vertices or edges 347
12.1 Introduction . . . 347
12.2 Proofs of the main results of this chapter . . . 350
12.3 Application: (bi)directed tree-width and cycle-width . . . 356

IV Coarse graph theory 361

13. Path-decompositions into parts of small radius 363
13.1 Introduction . . . 363
13.2 Preliminaries . . . 368
13.3 Graph-decompositions and quasi-isometries . . . 369
13.4 Radial path-width . . . 376
13.5 Radial cycle-width . . . 383
13.6 Radial star-width . . . 384

13.A Where quasi-geodesic topological minors are not the right obstruction . . . 391
13.B Counterexamples to stronger versions in the star and tree case 394

14. Characterising graphs quasi-isometric to K^4-minor-free graphs 399
14.1 Introduction 399
14.2 Preliminaries 402
14.3 Preparatory work: sequences of graph-decompositions 405
14.4 Proof of Theorem 38: Conjecture 36 for $X = K_4^-$ 409
14.5 Proof of Theorem 37: Conjecture 36 for $X = K_4$ 414

15. A Menger-type theorem for two induced paths 433
15.1 Introduction 433
15.2 Interlaced systems of intervals 435
15.3 Proof of Theorem 40 437

Bibliography 447

Appendix 467

A. Summary 469

B. Deutschsprachige Zusammenfassung 473

C. Publications related to this thesis 479

D. Declaration of contributions 481

E. Acknowledgements 485

F. Eidesstattliche Versicherung 487

Introduction

In this thesis, we study various types of connectivity in graphs through the lens of tangles and separations. We exhibit the global and local aspects of a graph's structure, study the connections of tangles with concrete substructures in a graph, find a tree-decomposition that displays the infinities of an infinite graph, and investigate the interplay of graph-theoretic (connectivity) notions with the large-scale geometry of a graph.

All this work builds on methods and concepts from the *Graph Minor Project* by Robertson and Seymour [RS83–04] as well as further related and subsequent research. The Graph Minor Theorem [RS04] of Robertson and Seymour confirms 'Wagner's conjecture', which asserts that in every infinite sequence of finite graphs, there exist two graphs such that one is a minor of the other (see [D25]). Their proof of the Graph Minor Theorem spans more than 20 papers and 500 pages [RS83–04]. In its course, Robertson and Seymour developed various tools and techniques that have heavily influenced the field of structural graph theory ever since.

One of their key innovations is the approach to view highly connected structures or 'clusters' in a graph through the lens of *tangles* [RS91]. Instead of describing clusters by what they consist of, i.e. the vertices or edges belonging to them, tangles capture clusters indirectly by a system of pointers to the location of the cluster in the graph.

The idea behind this approach can be described as follows: Intuitively, the majority of any highly connected structure or cluster lies, because of its high connectivity, on one side of each narrow bottleneck of the graph (formalised as low-order separations). The cluster therefore gives rise to a pointer towards the side of the bottleneck in which its majority is contained (formalised as oriented separations). In this way the cluster induces a whole system of pointers (formalised as an orientation of all of the graph's low-order separations) that points towards the cluster's location in the graph. Robertson and Seymour's innovation was to axiomatise the systems of pointers arising from clusters without any reference to the clusters themselves, but just in terms of the relation of the pointers to each other (formalised as 'consistent' orientations). This paved the way for studying the clusters of a graph with the unifying framework of tangles, which captures various kinds of concrete clusters at once.

The structure of a finite graph with respect to its tangles is tree-like in two different ways, both formalised in terms of special tree-decompositions and first established in [RS91]. On the one hand, the *Tree-of-Tangles Theorem* asserts that each finite graph admits a tree-decomposition that captures the relative position of the tangles with respect to each other in that distinct tangles are accommodated in distinct parts of the tree-decomposition. On the other hand, the *Tangle-Tree Duality Theorem* says that in the absence of tangles of high order (that is, large clusters), a finite graph admits a tree-decompositions into small parts, each too small to accommodate a high-order tangle.

In the last three decades, the notion of tangles has attracted much attention and found applications beyond the realm of graphs. In particular, tangle theory has been axiomatised in a more formal, abstract setting, with the aim of making the methods applicable in other discrete contexts such as matroids or data sets. Even for this abstract notion of tangles, analogues of the above-described central two tangle-theorems still hold [D18a; D18b; DEW19; DHL19; DO19; DO21; EK21; EKT21; EKT22b].

In this thesis we study connectivity in graphs and beyond, especially through the lens of tangles and tangle theory. Our contributions can be broadly grouped into four parts: In Part I, we use tangle-theory to display the local-global structure of a graph, using the novel concept of local coverings. In Part II, we investigate the relationship between tangles and concrete substructures (such as clusters) in graphs, contributing in particular to the problem of finding vertex sets that induce a given tangle. In Part III, we study how tree-decompositions interact with various notions of connectivity in graphs, especially constructing a tree-decomposition that displays the different types of infinities in infinite graphs and extending Menger's Theorem to bidirected graphs. In Part IV, we look at the large-scale geometry of graphs through tools and methods from connectivity theory.

In what follows we motivate each of these parts and their chapters and describe their main ideas and results. Each chapter also comes with its own introduction, to which we refer the reader for more details on context and motivation as well as formal definitions of the respective terminology.

Part I: Graph-decompositions: using tangles to capture local structure

Tangles capture the highly connected substructures or 'clusters' of a graph, and as the Tree-of-Tangles Theorem asserts, their relative position to each other is a tree-like arrangement. Such a tree-like arrangement, however, does not always comply with the intuitive perception of the graph's structure along its clusters, which may well incorporate cycles (see Figure 1). The main issue for this lies in the global nature of the connectivity captured by tangles: The paths witnessing the high connectivity of a substructure can divert arbitrarily from the substructure itself. In other words, the clusters captured by tangles, such as for example clique minors whose branch paths are subdivided many times, need not be 'local', in contrast to the intuitively perceived local building blocks of the graph's global structure.

To overcome this issue, we develop *local covers* of graphs [DJKK26]: They are topological coverings of a given graph (viewed as a 1-complex) that reflect the graph's r-local structure, that is, its $r/2$-balls, but unfold every further global structure in a tree-like way; here, $r > 0$ is an integer parameter describing the intended degree of locality. This gives rise to a new *local-global paradigm*: 'r-local in G corresponds to global in G_r'. For example, any notion of r-local separations should have the property that the r-local separations of G naturally correspond to the global separations of the r-local covering G_r.

With r-local coverings and the local-global paradigm, we offer a possible formal basis for studying the local and global aspects of graph properties. In Part I of this thesis, we formally develop this novel approach and use it to investigate the connectivity structure of

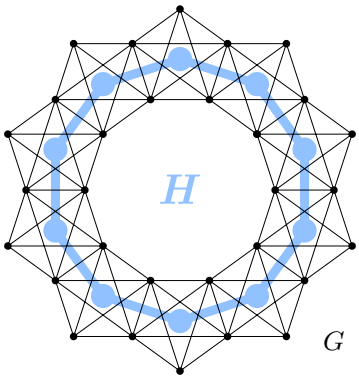

Figure 1: The global structure of G is displayed by a cycle H. Its local parts are K^5s.

graphs by applying concepts from tangle-theory. Our results have since found applications in the study of locally chordal graphs (see [AABKN25; AK25; AKK25] and Chapter 3) and in group theory [CKKT25; KS26].

In Chapter 1, we use this approach to capture the r-global structure of a graph G by the means of a suitable graph-decomposition $\mathcal{H}_r(G)$. In Chapter 2, we develop a theory of r-local separations compatible with r-local coverings and apply it to compute $\mathcal{H}_r(G)$ for many finite graphs G. In Chapter 3, we study canonical tree-decompositions of chordal graphs and use them to characterise r-locally chordal graphs in terms of their $\mathcal{H}_r(G)$.

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 $r > 0$ 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 decomposition $\mathcal{H}_r(G)$, and in particular the model, are obtained as projections of the tangle-tree structure of the r-local cover G_r of the given graph G. Recall that the r-global structure of a graph G is unfolded in a tree-like way in its r-local cover G_r. Such global tree-structure can be suitably captured in terms of the tangles of G_r. The Tree-of-Tangles Theorem, however, does not extend verbatim to infinite graphs; they need not admit a tree-decomposition that efficiently distinguishes all their tangles (see also Chapter 6). We prove the existence of such tree-decompositions for the usually infinite, but highly symmetric r-local cover G_r of a finite graph G, and we find the tree-decompositions even canonically, that is, such that they are invariant under the graph's automorphisms.

The canonicity of such a canonical tree-decomposition of G_r then allows us to 'fold' it along the r-local covering map $p_r : G_r \to G$, by taking the quotient under the natural action of the deck-transformation group of p_r on the tree-decomposition. This folding results in a canonical graph-decomposition of the graph G that decomposes G along its highly connected r-local parts and captures their relative position as the graph's r-global connectivity structure. In this way, the tangle theory from graph minors is brought to bear canonically on arbitrary graphs, which need not be 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 global structure is displayed by a graph.

Chapter 2: Canonical graph-decompositions via local separations

The r-local cover G_r of a finite graph G is usually infinite, so the above construction of the canonical graph-decomposition $\mathcal{H}_r(G)$ from Chapter 1 is often inherently infinite itself. This raises the question how we may compute $\mathcal{H}_r(G)$.

In this chapter, we find a way to reconstruct the decomposition $\mathcal{H}_r(G)$ from finite information in G itself that is sufficiently local to be reflected in the r-local cover G_r. For this, we develop a finite combinatorial theory of local separations, and we build on it to give a different construction of $\mathcal{H}_r(G)$, which allows us to compute $\mathcal{H}_r(G)$ for many finite graphs G. Our work extends the global theory of connectivity and separations to the local realm, and it generalises the notion of local 2-separations by Carmesin [C22].

With our definition of local separations, we achieve two seemingly opposite goals: On the one hand, local separations are defined purely in terms of local information in G, while on the other hand our definition follows the local-global paradigm in that the local separations of G naturally correspond to the global separations of the r-local covering G_r of G. This correspondence rests on a bounded computable graph invariant, the *r-displacement* of G, which needs to be large in terms of r and the order of the considered (local) separations (see Section 2.5). Under this assumption, our work yields the desired computability of $\mathcal{H}_r(G)$.

Chapter 3: Canonical tree-decompositions of chordal graphs

One of the natural questions arising from the definition of the graph invariant $\mathcal{H}_r(G)$ in Chapter 1 is: Can we characterise classes of graphs G in terms of properties of their $\mathcal{H}_r(G)$? In this chapter we give a first example of such a characterisation: We show that a locally finite, connected graph G is *r-locally chordal* (that is, its $r/2$-balls are chordal) if and only if the graph-decomposition $\mathcal{H}_r(G)$ is *into cliques* (that is, all its parts are complete subgraphs of G).

The proof of this characterisation relies on the study of locally finite chordal graphs in terms of their tree-decompositions. Halin characterised the chordal locally finite graphs as those that admit a tree-decomposition into cliques [H64; H89]. We show that these tree-decompositions can be chosen to be canonical, that is, so that they are invariant under all the graph's automorphisms.

From this, we deduce our characterisation of r-locally chordal graphs G in terms of $\mathcal{H}_r(G)$ along the local-global paradigm: It is easy to show that a graph is r-locally chordal if and only if its r-local cover is chordal [AKK25]. For a given r-locally chordal graph G, our results then imply that the canonical tree-decomposition of the r-local cover G_r that defines $\mathcal{H}_r(G)$ by folding along p_r is into cliques. This directly implies that if G is r-locally chordal, then $\mathcal{H}_r(G)$ is into cliques. The converse implication also holds, but needs further observations about tree-decompositions and about coverings and cliques.

Part II: Tangles and concrete clusters

Tangles are designed as an indirect way of capturing clusters in graphs by *consistent* orientations of low-order separations of graphs [RS91]. Robertson and Seymour distilled the notion of consistency from the orientations arising via majority vote by the clusters of a graph: for each low-order separation, the high connectivity of a cluster implies that a majority of the cluster lies on one side of the separation, towards which we then orient the separation. The abstract notion of tangles captures all such orientations induced by clusters via majority vote. This naturally raises the question whether the converse also holds, that is, whether every tangle in a graph is induced by a cluster via majority vote.

Qualitatively, this question has an affirmative answer in the following sense: there is a function $f : \mathbb{N} \to \mathbb{N}$ such that for every k-tangle τ in a graph G, there is a grid minor in G of size $f(k) \times f(k)$ which induces the subtangle of τ of order $f(k) < k$ [RS91]. If we look for an exact converse, however, the question is still open even in the following weak form, first proposed by Diestel, Hundertmark and Lemanczyk [DHL19]:

Question 1. *Is every k-tangle τ in a graph induced by some vertex set X by majority vote in that $|X \cap A| < |X \cap B|$ for every separation $(A, B) \in \tau$?*

Substantial progress towards Question 1 was made by Elbracht, Kneip and Teegen [EKT20], who give an affirmative answer to the following fractional relaxation of Question 1: for the majority vote, we allow the vertex set V of the graph to be equipped by not just an assignment of values in $\{0, 1\}$ (giving a vertex subset), but by a more general weight function $w : V \to \mathbb{Z}_{\geqslant 0}$ assigning non-negative integers to the vertices. Moreover, Question 1 is known to hold for tangles of small order $\leqslant 3$ [E15; G16; RS91], and it should be possible to derive the case $k = 4$ from the recent characterisation of 4-tangles by Carmesin and Kurkofka [CK25].

We make two contributions towards a resolution of Question 1: In Chapter 4, we reduce Question 1 to graphs of size bounded only in k, using a new inductive proof method for k-tangles. In Chapter 5, we study Question 1 for tangles in discrete contexts more general than graphs and provide two characterisations for the existence of point sets X inducing such tangles.

In Chapter 6, we find a concrete substructure of a graph not from a single tangle, but from the graph's tangle-structure: We show that Robertson and Seymour's Tree-of-Tangles Theorem extends to locally finite graphs without thick ends, by exhibiting thick ends (and hence half-grid minors) at every obstruction to the existence of tangle-distinguishing tree-decompositions.

Chapter 4: On vertex sets inducing tangles

We reduce Question 1 for k-tangles to graphs whose size is bounded by a function in k. This in particular implies that for any given integer k, the question whether Question 1 holds for k-tangles can be computationally decided in finite time.

Additionally, we show that if Question 1 has a positive answer for some fixed k, then every k-tangle is induced by a vertex set whose size is bounded in k. A similar result holds for the fractional relaxation of Question 1: Building on the existence of tangle-inducing

weight-functions shown by Elbracht, Kneip and Teegen [EKT20], we show for all k that every k-tangle in a graph G is induced by a weight function $V(G) \to \mathbb{Z}_{\geqslant 0}$ whose total weight is bounded in k.

As the key step of our proofs, we show that any given k-tangle in a graph G is the lift of a k-tangle in some topological minor of G whose size is bounded in k. This provides a new inductive proof method for tangles in terms of the size of their host graphs, which we hope to be of independent interest. In order to find such a topological minor for a given k-tangle in G, we either make use of a tangle of order at least $k+1$ or, in the absence of such tangles, we build on the Tangle-Tree Duality Theorem to carefully decompose the graph in a highly regular way which allows us to control the separations that arise by deleting a suitably chosen edge of the graph. The proof spans more than 25 pages and is the technically most complex and challenging one in Part II of this thesis.

Chapter 5: Point sets and functions inducing tangles of set separations

In general discrete structures, the analogue of Question 1 has a negative answer: not all tangles are induced by a set of points (see [EKT20]). We aim to characterise those tangles that are: the tangles that are induced by a subset of or a function on the set of data points whose connectivity structure they are meant to capture.

We offer two such characterisations. The first is in terms of how many small sides of a tangle's separations it takes to cover the ground set. The second uses a new notion of duality for oriented set separations that becomes possible if these are no longer required to be separations of a graph or matroid.

Chapter 6: Efficiently distinguishing all tangles in locally finite graphs

Robertson and Seymour's Tree-of-Tangles Theorem, one of their central two tangle theorems, asserts that each finite graph has a tree-decomposition that efficiently distinguishes all its tangles [RS91]. This, however, does not extend verbatim to infinite graphs: there are even locally finite, planar graphs that do not admit such tree-decompositions [AJKP25; EKT22b].

We prove that thick ends are the only obstruction that may prevent a locally finite graph from admitting a canonical tree-decomposition that efficiently distinguishes all its tangles. More precisely, we exhibit a thick end (and hence a half-grid minor) at any obstruction to the existence of such tree-decompositions. Our proof builds on new methods for the analysis of the limit behaviour of strictly increasing sequences of separations.

Part III: Connectivity and tree-structures

A recurring topic in structural graph theory are dualities such as the Tangle-Tree-Duality Theorem; it asserts that the existence of a high-order tangle is dual to the existence of a tree-decomposition with small parts: the absence of high-order tangles can be witnessed by a tree-decomposition whose parts are too small to accommodate high-order tangles (see [A24b; DO19; DO21; RS91]). Such tree-decompositions have small *tree-width*, which is defined as the maximum over the cardinality of the tree-decomposition's parts (minus 1).

Structural dualities like the Tangle-Tree Duality exist for many different width-measures (like tree-width) and obstructions (like high-order tangles); the perhaps most classical examples are the (directed) grid theorems for (directed) tree-width [KK15; RS86b]. Structural dualities form an important method to understand the type of connectivity captured by a given width-measure and its corresponding obstructions as well as a suitable tool to efficiently solve computing problems for graphs of small such width (see e.g. [B06]). Both these applications fuel the quest to find strongest possible dualities in terms of restrictive width-witnessing (tree-)decompositions and of small or simple lists of obstructions.

In Chapter 7, we find a tree-decomposition displaying the infinities of graphs of finite tree-width, by generalising Thomas and Kříž' result on linked tree-decompositions of graphs of finitely bounded tree-width [KT91]. In Chapter 8, we provide counterexamples to various possible generalisations of the main result of the previous chapter. In Chapter 9, we prove a structural duality for strong immersions of walls in terms of special tree-cut decompositions. In Chapter 10, we extend the Lovász-Cherkassky theorem from locally finite graphs to locally finite graphs with ends. In Chapter 11, we investigate whether and how Menger's theorem generalises to bidirected graphs. In Chapter 12, we study hitting sets for cycles through prescribed sets and vertices, and we derive improved bounds on (bi)directed tree-width and cycle-width.

Chapter 7: Linked tree-decompositions into finite parts

A classical result of Kříž and Thomas asserts that every graph of finitely bounded tree-width has a (rooted) tree-decomposition which witnesses the graph's tree-width and which is additionally *linked* [KT91]. This result forms an important step in the work of Robertson and Seymour on the Graph Minor Theorem (see [RS90a]) and its extension to infinite graphs of finitely bounded tree-width by Thomas [T89]. Roughly speaking, a tree-decomposition is *linked* if it captures all connectivity in the graph along its (rooted) decomposition tree.

We extend the result of Kříž and Thomas to graphs which admit a tree-decomposition into finite parts, including those of finite, but not finitely bounded tree-width: we show that such a graph also has a rooted tree-decomposition into finite parts that is linked, tight and componental. This result is the most profound and far-reaching result in Part III of the thesis, and its proof requires a deep understanding of the interaction between tree-decompositions and the specifics of connectivity in infinite graphs. We remark that while achieving just a subset of 'linked', 'tight' and 'componental' may be significantly easier, achieving all three of them simultaneously poses a significant challenge. It is their surprisingly powerful interplay which enables several applications.

As a first application, we obtain that every graph without half-grid minor has a lean tree-decomposition into finite parts, extending the corresponding result by Kříž and Thomas for graphs of finitely bounded tree-width [KT91]. In particular, it follows that every graph without half-grid minor has a tree-decomposition which efficiently distinguishes all its ends and critical vertex sets, strengthening results by Carmesin [C19] and by Elm and Kurkofka [EK22] for this graph class.

As a second application of our main result, it follows that every graph G which admits a tree-decomposition into finite parts has a tree-decomposition into finite parts that displays

all infinities of G, that is, it displays the ends of G homeomorphically, their combined degree and their dominating vertices (resolving a question of Halin from 1977 [H77]) as well as the critical vertex sets of G and their tight components cofinitely. From this tree-decomposition, we derive as direct corollaries unified proofs of the characterisations of graphs without half-grid minor [RST95] and of graphs without binary-tree subdivision [ST93] due to Robertson, Seymour and Thomas. Moreover, Albrechtsen applied this tree-decomposition to extend the Tangle-Tree Duality Theorem of Robertson and Seymour to infinite graphs [A24b].

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

We provide counterexamples to various possible generalisations and strengthenings of the results in Chapter 7, demonstrating that our work achieves very much the best one can hope for in this setting. In particular, we construct a locally finite, planar, connected graph that has no lean tree-decomposition.

Chapter 9: A grid theorem for strong immersions of walls

The grid theorem of Robertson and Seymour asserts that a graph either has small tree-width or contains a large wall as a (topological) minor [RS86b]. This classical structural duality has analogues in various settings such as directed graphs [KK15] or for other containment-relations such as weak immersions of walls [W15].

Weak and strong immersions are generalisations of topological minors, where the paths connecting the branch vertices are allowed to cross; for strong immersions, the internal vertices of a path may not be branch vertices themselves. Dvořák and Wollan proved a structure theorem for graphs excluding a strong immersion of a fixed graph [DW16], but this theorem admits no converse for walls. We give a precise structural duality: a graph contains no large wall as a strong immersion minor if and only if the graph admits a tree-cut decomposition of small 'width', which is measured in terms of its adhesion and the path-likeness of its torsos.

Chapter 10: The Lovász-Cherkassky theorem for locally finite graphs with ends

For a set T of vertices of a finite graph G, let $\lambda_G(t, T \setminus \{t\})$ be the size of a smallest cut that separates t and $T \setminus \{t\}$. Clearly, there are at most $\lambda_G(t, T \setminus \{t\})$ edge-disjoint paths connecting t and $T \setminus \{t\}$, and hence every family of edge-disjoint T-paths (i.e. the paths are internally disjoint from T and connect distinct vertices of T) has size at most $\frac{1}{2}\sum_{t\in T} \lambda_G(t, T \setminus \{t\})$. Lovász [L76] and Cherkassky [C77] discovered independently that if each vertex $v \in V(G) \setminus T$ has even degree, then this upper bound is attained. In other words, in any maximum-sized family of edge-disjoint T-paths, there are $\lambda_G(t, T \setminus \{t\})$ many paths between t and $T \setminus \{t\}$ for every vertex $t \in T$.

Based on this structural reformulation, Jóo extended the theorem of Lovász and Cherkassky to locally-finite infinite graphs [J23]. Diestel asked whether one can even extend the theorem to locally finite graphs with their ends [D20]. We answer this question in the affirmative. In our result, T may contain not just vertices but also ends, and paths are one-way (two-way) infinite when they establish a vertex-end (end-end) connection.

Chapter 11: Menger's theorem in bidirected graphs

Bidirected graphs are a generalisation of (un)directed graphs, in which every edge is assigned a sign + or − at each of its endvertices; they were first introduced by Kotzig in a series of papers [K59a; K59b; K60]. Recently, bidirected graphs have received increasing attention both from an algorithmic perspective [AFN96; BZ06] and in structural graph theory [AK06; W22], where they are of particular interest due to their correspondence to undirected graphs with perfect matchings and Norine's matching grid conjecture [N05]; see [W22].

The connectivity structure of bidirected graphs is considerably more complex than the one of (un)directed graphs and far less understood. One particular complication is that the existence of a walk between two vertices does not imply the existence of a path joining them. We here investigate whether and how one of the basic cornerstones of connectivity theory, Menger's theorem [M27] (the minimum size of a vertex set separating two sets X and Y of vertices equals the maximum number of disjoint paths connecting X and Y), extends to bidirected graphs. We show that Menger's theorem is generally not true in bidirected graphs. Nonetheless, we exhibit a sufficient condition for X and Y which yields a version of Menger's Theorem in bidirected graphs that in particular extends its (un)directed analogues.

Chapter 12: Hitting cycles through prescribed vertices or edges

Hatzel, Rabinovich and Wiederrecht [HRW19] introduced the cycle-width of (bi)directed graphs as a tool to transfer results about (bi)directed tree-width to the respective corresponding undirected graphs with perfect matchings. We here show that the (bi)directed tree-width of a directed graph is linearly bounded in its cycle-width, which improves the previous quadratic upper bound by Giannopoulou, Kreutzer and Wiederrecht [GKW24] with a much simpler proof.

Surprisingly, we obtain this result by answering a natural question about hitting sets for cycles through prescribed sets of vertices or edges: We prove that for every set S of vertices of a directed graph D, the maximum number of vertices in S contained in a collection of vertex-disjoint cycles in D is at least the minimum size of a set of vertices that hits all cycles containing a vertex of S. We also show that the corresponding statement in bidirected graphs is true and that its edge-variant holds in both undirected and directed graphs, but fails in bidirected graphs. The vertex-version in undirected graphs remains an open problem.

Part IV: Coarse graph theory

A recent strand of graph theory views graphs from a metric perspective and aims to develop a 'coarse' graph theory to understand the large-scale geometry of graphs. The term 'coarse graph theory' was introduced by Georgakopoulos and Papasoglu [GP25], who followed Gromov's approach to coarse geometry [G93] into the realm of graphs and presented a collection of problems and results on the interplay of graph minors and geometry. A central notion to Gromov's approach is that of a *quasi-isometry* between metric spaces, that is, a

bi-Lipschitz map that allows for an additive error. This notion also becomes the central tool of Georgakopoulos' and Papasoglu's approach: Quasi-isometries provide the formal framework to study the coarse geometry of graphs.

Independently from their research, our work was initiated by the use of graph-decompositions in Part I. We set out to find a width-measure which measures how much the large-scale geometries of the decomposed graph and the decomposition graph coincide. A first natural attempt would be to extend the notion of tree-width verbatim, that is, the width of a decomposition is defined as the maximum over the cardinalities of its parts (minus 1). However, as Diestel and Kühn observed [DK05], this yields no meaningful extension of tree-width (see Section 13.1.2 for more details).

This led us to take a metric perspective instead: With the notion of *radial width*, we measure the width of a graph-decomposition as the maximum over the radius of its parts, not their cardinality. It turns out that if a graph-decomposition is 'honest' and has small radial width and also small 'radial spread', then this is, in fact, equivalent to the existence of a bounded-parameter quasi-isometry between the decomposed graph and the decomposition graph (see Section 13.1.3 for the details). Our viewpoint via graph-decompositions has the advantage of providing a concise framework for the explicit construction of a quasi-isometry between the decomposed graph and the decomposition graph.

We then set out to identify obstructions to the existence of graph decompositions of small radial width whose decomposition graph H belongs to some prescribed (minor-closed) class $\mathcal{H}$ of graphs (such as the class of trees for tree-decompositions). In analogy to the theory of graph minors, this required a coarse, metric version of minors and topological minors. We ended up with notions very similar to Georgakopoulos and Papasoglu, who established the so-called 'fat minor conjecture':

Conjecture 1 ([GP25, 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, describing that the minor is still seen from 'further away' (parametrized in K). 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 13.3.4 for the formal definition).

By now Conjecture 1 is known to hold for several small cases; see [ADG25; CDEHV08; FP23; GP25] and Chapters 13 and 14. Our proofs build on graph-decomposition based methods, and Chapter 14 also uses a coarse version of Menger's theorem for two paths (see Conjecture 2), which we prove in Chapter 15. In contrast, Conjecture 1 is false for general sets $\mathcal{X}$ [DHIM24], even in a substantially weaker form [AD25], and also for small graphs in $\mathcal{X}$ [ADG26]. Together with the aforementioned positive results, this leaves $\mathcal{X} = \{K^5\}$ and $\mathcal{X} = \{K_{3,t}\}$ for $t \geqslant 3$ as the critical open cases; in particular, $\mathcal{X} = \{K^5, K_{3,3}\}$ remains open, which captures quasi-isometry to planar graphs and is therefore also known as the 'Coarse Kuratowski Conjecture' [GP25, Conjecture 9.1].

Chapter 13: A structural duality for path-decompositions into parts of small radius

We prove Conjecture 1 for $\mathcal{X} = \{K^3, K_{1,3}\}$ (quasi-isometric to a disjoint union of paths), $\mathcal{X} = \{K_{1,3}\}$ (cycles and paths) and $\mathcal{X} = \{K^3, W\}$ (subdivided stars), where W is the 'wrench graph' obtained from two paths of length 2 by joining their middle vertices with an edge (see Figure 13.1). In all these cases, Conjecture 1 even holds true when replacing 'fat minors' by the stronger notion of 'quasi-geodesic topological minors', which can be thought of as a coarse analogue of topological minors.

Chapter 14: A characterisation of the graphs quasi-isometric to K^4-minor-free graphs

We confirm Conjecture 1 for $\mathcal{X} = \{K^4\}$. To date this is one of the most challenging cases in which Conjecture 1 is known to have an affirmative answer. Our proof builds on the correspondence between quasi-isometries and graph-decompositions of small radial width and radial spread. We construct such a graph-decomposition iteratively via a sequence of decompositions of larger and larger subgraphs until we finally decompose the whole graph (see Section 14.3). In each iteration step, we find either a fat K^4 minor or a suitable larger subgraph to which we can extend our decomposition without constructing a K^4 minor in the extended decomposition graph. Our proof technique also yields a new short proof of the respective K_4^--case, which was first established by Fujiwara and Papasoglu [FP23].

Chapter 15: A Menger-type theorem for two induced paths

In order to verify Conjecture 1 for a set $\mathcal{X}$ of graphs with our graph-decomposition based method, we often work iteratively and extend the domain of a partial decomposition until we have decomposed the whole graph – unless we find a fat X minor for some $X \in \mathcal{X}$ in some iteration step. An often helpful tool to find X minors is Menger's theorem, and for our work in the coarse realm, we conjectured the following coarse Menger-type theorem:

Conjecture 2. *There is a function $f : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that for all $k, d \in \mathbb{N}$ and for all graphs G and vertex sets $X, Y \subseteq V(G)$, either there exist k disjoint X–Y paths, pairwise at distance at least d, or there exists a set $Z \subseteq V(G)$ of size at most $k-1$ such that each X–Y path has distance at most $f(k,d)$ to some vertex in Z.*

In this chapter, we confirm Conjecture 2 for $k = 2$ with $f(2,d) := 129d$; this case serves as in important tool in Chapter 14. Independently from us, Georgaokpoulos and Papasoglu also proposed Conjecture 2 and proved it for $k = 2$ with the slightly worse function $f(2,d) := 272d$. A further proof was later given by Nguyen, Scott and Seymour in [NSS25a].

Conjecture 2 and our result immediately gained a lot of attention. While Conjecture 2 is now known to hold in various restricted settings [GKL23; GHW25; HNST24; NSS25d; NSS25e], Nguyen, Scott and Seymour [NSS25a] disproved Conjecture 2 for all $k \geqslant 3$ and $d \geqslant 3$, even in the weaker form where the size of Z is also allowed to depend on k and d [NSS25c]; the case $d = 2$ remains open for all $k \geqslant 3$.

PART I

Graph-decompositions: using tangles to capture local structure

CHAPTER 1

Canonical graph-decompositions via coverings

This chapter is based on [DJKK25; DJKK26] and joint work with Reinhard Diestel, Paul Knappe and Jan Kurkofka.

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 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.

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 [S22]. 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.

Our main result reads as follows:

Theorem 3. *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.*

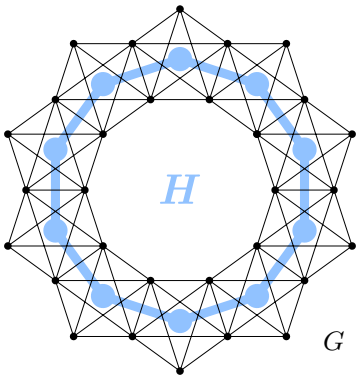

Figure 1.1: The global structure of G is displayed by a cycle H. Its local parts are K^5s.

See Sections 1.1.1 to 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 3 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 4. *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 4. Similarly, we shall see that the r-local structure of any finite graph G that complements its r-global structure as described in Theorem 3 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 [CKKT25]. More group-theoretical applications of this approach have been obtained by Carmesin [C].

The idea for our construction of $H = H(G, r)$ was inspired by the aims of Carmesin [C22]. It is to apply the theory of tangles from graph minor theory 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. 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 3 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 [S72], 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 3), but also when it is the Cayley graph of an arbitrary finitely generated group (Theorem 4). 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 3 precise, and to give a brief sketch of its proof.

1.1.1 Tree-decompositions vs. graph-decompositions

Early in their seminal work on graph minors, Robertson and Seymour [RS83–04] 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 3 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 [D25].

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 witness Theorem 3, 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 3 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; EKT22b]. 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 $\mathcal{D}_r$ 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 $T_r/\mathcal{D}_r$, 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 $\mathcal{D}_r$-orbit of t.

1.1.3 Details of Theorem 3

In order to make Theorem 3 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 $\mathcal{D}_r$ 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 [RS83–04]. 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 $\mathcal{D}_r$). 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 $\mathcal{D}_r$ (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 $N/\mathcal{D}_r$. This is not a graph. But we show that it can be displayed by a graph, the graph H from Theorem 3, 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 [D18a]. 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 $H = T_r/\mathcal{D}_r$ 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 $\mathcal{D}_r$-orbit of t, *display* the r-global structure $(N, \mathcal{D}_r)$ 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 Theorem 3 and Theorem 4

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 [H18] 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 [W71], by Stallings's group splitting theorem [S72]. 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 5. *Accessible locally finite, quasi-transitive graphs are in fact tangle-accessible.*

Applied to Cayley graphs of groups, Theorem 5 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 3, 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 3. 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 3 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 3 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 [JKK25] 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 [C22] 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 the 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 3. The proof is then completed in Sections 1.6 to 1.8.

Acknowledgement

A referee of our paper [DJKK26] on which this chapter is based suggested that we point out the following more topological way to define the r-local cover of G. Rather than taking the covering space of G with $\pi_1^r(G, v_0)$ as characteristic subgroup, we could glue a disc on every cycle of length at most r in G and then work with the universal cover of the resulting 2-complex. We are happy to invite the reader to keep this alternative perspective in mind as the chapter proceeds.

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 [D25], 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 [L17b]. 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 [H02]. 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 [D25], and say *topological path* for continuous maps $[0, 1] \to G$.

In addition to Hatcher's [H02, Ch.1.3] we use standard covering space terminology as in Jänich [J84, 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 $\mathcal{D}(p)$. As generally with automorphisms, we write multiplication in $\mathcal{D}(p)$ 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 [S83], on which our terminology is based.

We define a *walk* (of *length* k) in a graph G to be 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 [S83]. 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 [D25, Lemma 12.3.1].

A precursor to our notion of graph-decompositions given below 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 a *graph-decomposition with model* H, or *H-decomposition*, of G 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.

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,[1] the maximum of its co-bag sizes $|W_v|$ its *co-width*, or *spread*. 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 $(H, \mathcal{V})$.

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

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.

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 $V_t \cap V_{t'}$ [D25, 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 $S := \big(\bigcup_{h \in U \cap W} V_h\big) \cup \bigcup_{hh' \in F} (V_h \cap V_{h'})$, 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. Indeed, for adjacent $v, v' \in f(X_Y)$ consider any $h \in V(H)$ such that $v, v' \in V_h$, as given by (H1). Then $h \in W_v \cap W_{v'}$. In particular, $W_v \cap W_{v'} \neq \emptyset$, so the components of $H - Y$ that contain W_v and $W_{v'}$, respectively, are identical. As $f(X_Y)$ is connected, this implies that 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 [D25] 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 $(H, \mathcal{V})$ 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 $(H, \mathcal{V})$ is honest; 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 $t_1 T t_2$ or $t_3 T t_0$, 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 $\mathcal{D} \subseteq \mathrm{Aut}(C)$ be the group of deck transformations, and let $\mathcal{T} = (T,\mathcal{V}_T)$ be any regular $\mathcal{D}$-canonical tree-decomposition of C. By Lemma 1.3.7, every $\varphi \in \mathcal{D}$ 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 $\mathcal{D}$ on T, since for every $\varphi \in \mathcal{D}$ 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 $T/\mathcal{D}$ of the action of $\mathcal{D}$ 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 $\mathcal{D}$-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, $\mathcal{D}$ 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 $C[T_t]$ 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 $\mathcal{D}$-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 $H = T/\mathcal{D}$, 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 3 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 ([D74, 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. Since $\mathcal{T}$ is canonical, $\hat{\varphi}$ acts on the decomposition tree T, uniquely by Lemma 1.3.7, as of course does $\mathcal{D}$. Since $\hat{\varphi}$ is a lift of φ, its action on T is well defined on the orbits of $\mathcal{D}$ in T, i.e., $\psi\colon \mathcal{D} \cdot t \mapsto \mathcal{D} \cdot \hat{\varphi}(t)$ is a well-defined automorphism of $H = T/\mathcal{D}$. 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 $H = T/\mathcal{D}$ 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 $\mathcal{D}(p)$-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 so does $(H, \mathcal{G})$. If $(H, \mathcal{G})$ has finite spread and G is 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.

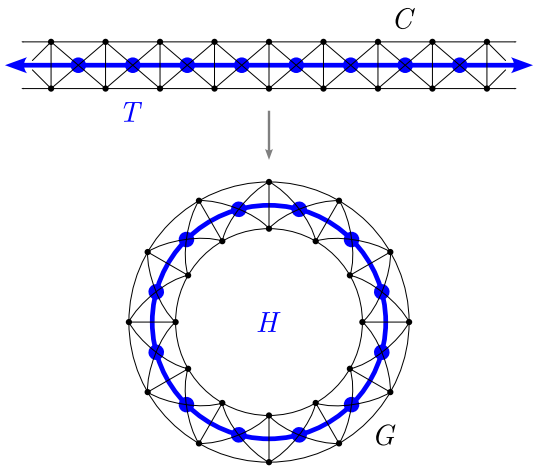

Figure 1.2: The global structure of G is displayed by a cycle. Its local parts are K^4s.

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 $\mathcal{D}$ 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 model $H = T/\mathcal{D}$ 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 $\mathcal{D}$ of this cover is the infinite cyclic group $\mathbb{Z}$, which acts on C by shifting.

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 model $H = T/\mathcal{D}$ 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. □

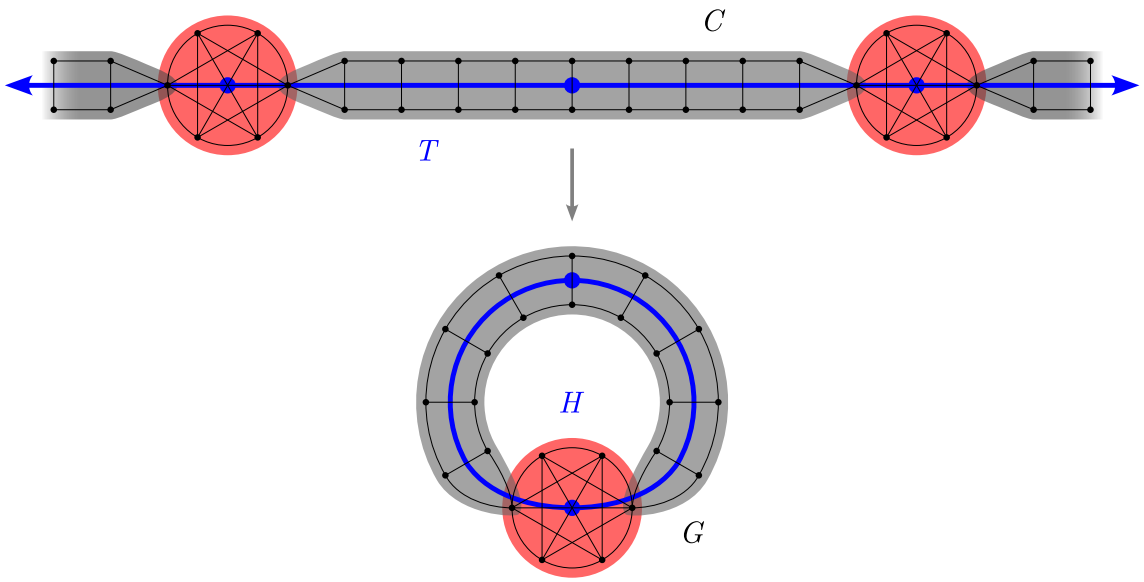

Figure 1.3: The grey part of H is not an induced subgraph of G

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 model H of G. Indeed, one can show that T is a cover of H if and only if the deck transformation group $\mathcal{D}(p)$ 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. [H02, 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 $\mathcal{D}$ 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 model, 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 $\mathcal{D}(p)$, 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 assertions hold:

(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 [H02] 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 $x_0Tuv Tx_0$ and its reverse $x_0Tvu Tx_0$ 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 $B_G(v, \varrho/2)$.[2] 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 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. □

[2] A notion that could be regarded as a precursor of the preservation of $(\varrho/2)$-balls has been studied by Georgakopoulos [G17].

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.

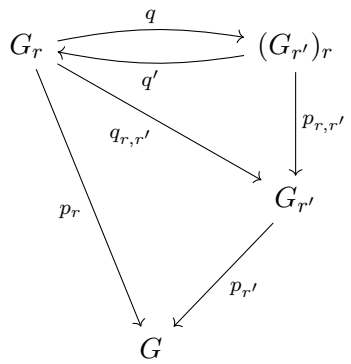

Figure 1.4: The interaction of local coverings for different values of r

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.*

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. □

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

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 [JKK25].

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 [B21] exhibits finite graphs with n-sheeted coverings for

any $n \in \mathbb{N}$.

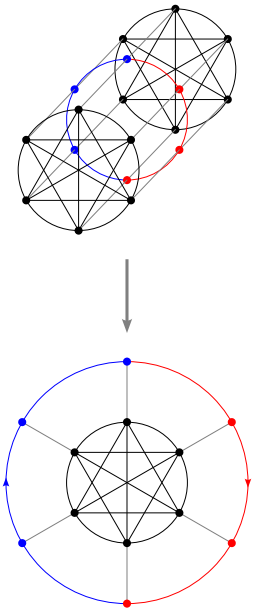

Figure 1.5: A graph whose 4-local covering is finite.

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 [M68; Š55] 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 $\mathcal{D}$ of all deck transformations of p has a finite generating set S such that C has a graph-decomposition modelled on $D := \mathrm{Cay}(\mathcal{D}, S)$ 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 $\varphi \in \mathcal{D}$, can be chosen as $C_\varphi = B_C(\varphi(\hat{x}_0), |G|)$.*

With these C_φ the graph-decomposition $(D, (C_\varphi)_{\varphi\in\mathcal{D}})$ 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 $\mathcal{D}$. If $\mathcal{D}$ is an accessible group, then C is an accessible graph.

Proof. Choose $(C_\varphi)_{\varphi\in\mathcal{D}}$ as in the statement of the theorem. As these C_φ depend only on $\mathcal{D}$, not on the choice of generators for our Cayley graph $\mathrm{Cay}(\mathcal{D}, S)$, 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 $\varphi \in \mathcal{D}$ 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 $\mathcal{D}$ such that, for $D = \mathrm{Cay}(\mathcal{D}, S)$, the pair $(D, (C_\varphi)_{\varphi\in\mathcal{D}})$ 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 $\{\varphi \in \mathcal{D} \colon \hat{v} \in C_\varphi\}$ 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 $\mathrm{Cay}(\mathcal{D}, S)$, and hence in $D_{\hat{v}_0}$.

The pairs of nodes of $D_{\hat{v}_0}$ are the pairs $(\varphi, \varphi') \in \mathcal{D}^2$ 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

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

and let $S := \bigcup_{v\in V(G)} S_v$. Then the subgraphs $D_{\hat{v}_0}$ of $D = \mathrm{Cay}(\mathcal{D}, S)$ 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 $\mathcal{D}$ acts on D by left-multiplication. Since our covering p is normal, there exists $\psi \in \mathcal{D}$ 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 $\chi \in \mathcal{D}$ 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 $(D, (C_\varphi)_{\varphi\in\mathcal{D}})$ is a graph-decomposition.

For a proof that S generates $\mathcal{D}$, we have to show that $D = \mathrm{Cay}(\mathcal{D}, S)$ 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 $(D, (C_\varphi)_{\varphi\in\mathcal{D}})$ 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 $D = \mathrm{Cay}(\mathcal{D}, S)$ 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 $\varphi_0, \varphi_0' \in \mathcal{D}$ 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 $\mathcal{D}$ 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 $(D, (C_\varphi)_{\varphi\in\mathcal{D}})$ 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 $\varphi \in \mathcal{D}$: that $(D, (C_\varphi)_{\varphi\in\mathcal{D}})$ has finite spread. By Lemma 1.3.5, therefore, the directions of C and of $D = \mathrm{Cay}(\mathcal{D}, S)$ are in bijective correspondence, and hence so are their ends [DK03, Theorem 2.2]. The ends of $\mathcal{D}$, 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 $\mathcal{D}$ 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 $\mathcal{D}_r$ 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 $\mathcal{D}_r = \pi_1(G, x_0)/\pi_1^r(G, x_0)$ 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 [D85]. 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 $N = \langle R \rangle_S^{\lhd}$ 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 $\mathcal{D}$ of all deck transformations of p is isomorphic to $\pi_1(G)/N$ and hence to Γ. Thus, $\mathcal{D}$ is inaccessible, and hence so are all its locally finite Cayley graphs [TW93].

Now consider the graph-decomposition $(D, (C_\varphi)_{\varphi\in\mathcal{D}})$ of C modelled on a locally finite Cayley graph D of $\mathcal{D}$ 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 $W = g_0e_0g_1 \dots g_{k-1}e_{k-1}g_k$ 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 [D85], 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 3 and 4 precise and give an overview of their proof. This will include a proof of Theorem 5.

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 3, 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 3 and 4, 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 3. This includes a common generalisation, Theorem 1.5.5, of Theorems 3 and 4. The outline will serve as our road map when we prove Theorems 3 and 4 (by way of proving Theorem 1.5.5) in Sections 1.6 to 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 [D25, 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 $S_{\aleph_0}$ [D17a]. 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 [D25, 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 [EKT22b, 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 3 and 4 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 $\mathcal{D}_r = \pi_1(G)/\pi_1^r(G)$ 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 $\mathcal{D}_r$ 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 $(N(G_r), \mathcal{D}_r)$, where $\mathcal{D}_r$ 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 3, or a Cayley graph as in Theorem 4.

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 $H = T/\mathcal{D}$. Its edges arise from ($\mathcal{D}$-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 $H_1 \preccurlyeq H_2 \dots$ 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 [JKK25].

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 [C22] 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 [AKK25] 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 3, for which we now give an outline. We shall prove Theorem 3 by way of proving a stronger result, Theorem 1.5.5, which will include Theorem 4 at no extra cost. A proof of Theorem 5 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 [EKT22b] (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 $\vec{N}'(G_r)$ 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 $\mathcal{D}_r$ 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 consider any infinite sequence $(A_0, B_0) < (A_1, B_1) < \dots$ in $\vec{N}'(G_r)$; we have to show $\bigcap_{i\in\mathbb{N}} B_i = \emptyset$. Since N'_{end} is exhaustive, it suffices to find in $\vec{N}'_{\text{end}}$ 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 $\mathcal{D}_r$, 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 3 and 4 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 3, are trivially quasi-transitive, and so are locally finite Cayley graphs such as the $\mathrm{Cay}(\Gamma, S)$ in Theorem 4.

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 3 and 4, 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 ([H18, 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 3

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 $N_{\text{end}} \subseteq N$ 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.[3]

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) [EKT22b, 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 $\overrightarrow{N}$ [EKT22b, 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 *ex-*

[3] In general, when N is infinite, the elements of $\overrightarrow{N}$ need not lie in a splitting star of N; see [D18a; GK21].

haustive 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:

Lemma 1.7.2 ([EKT22b, 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.*[4]

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 [EKT22b], 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 ([EKT22b])**.** *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.*

[4]The last sentence, which we already verified, is not explicitly stated in [EKT22b].

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 [D25, 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 $\overrightarrow{N}$ 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 $\overrightarrow{N}_{\text{end}}$, it will also be exhaustive (Lemma 1.7.3). Essentially, this amounts to showing that our sequence is interlaced with a sequence in $\overrightarrow{N}_{\text{end}}$: that at most finitely many (A_i, B_i) have the same $\leqslant$-relationships to all the elements of $\overrightarrow{N}_{\text{end}}$. 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.A below.

Recall that our plan is to show that our sequence $(A_i, B_i)_{i \in \mathbb{N}}$ in $\vec{N}$ dominates a suitable sequence in $\vec{N}_{\text{end}}$. 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 [D25, 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 $\overrightarrow{N}$, as follows. For every splitting star σ of N_{end} let

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

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

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

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

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. Let $r_i \in N_{\text{end}}$ be a separation oriented differently by O_i and O_{i+1}. Then one of these, $\vec{r_i}$ say, points towards $\vec{s_{i+1}}$, 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 $\vec{r'} \in \sigma$ points towards $g \cdot s$: that $s' := g \cdot s$ has an orientation $\vec{s'} \geqslant \vec{r'}$. Let $\vec{r} := g^{-1} \cdot \vec{r'}$. 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 $\vec{r'} = g \cdot \vec{r} \leqslant g \cdot \vec{s} =: \vec{s'}$ 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 3 that also includes Theorem 4 – 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 [D25, 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 [D25, 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.

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 3 and 4, as outlined at the start of this section. More precisely, what we shall need is Theorem 5, which we will thus prove first. Our input N for Lemma 1.8.1 in the proof of Theorem 5 will come from the following non-trivial result:

Lemma 1.8.4 ([EKT22b, Theorem 6.6])**.** *Every connected, locally finite graph has a canonical tree of tangles consisting of relevant separations.*[5]

We can now prove Theorem 5, which we restate:

Theorem 5. *Accessible locally finite, quasi-transitive graphs are even tangle-accessible.*

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. □

[5]The relevance of the separations is not explicitly stated in [EKT22b, Theorem 6.6] but immediate from its proof.

Note that the tree of tangles $N'(G)$ given by Lemma 1.8.4, which we just used in the proof of Theorem 5, 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 5, 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 5 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 3. Theorem 3 is the special case of Theorem 1.5.5 where G is finite. □

Proof of Theorem 4. Since $\mathrm{Cay}(\Gamma, S)$ is locally finite and transitive, Theorem 1.5.5 implies Theorem 4. □

1.A Tree-decompositions in quasi-transitive graphs

In this appendix 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.A.1 to 1.A.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.A.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.A.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.A.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.A.3: that the parts of $\mathcal{T}$ have at most one end:

Lemma 1.A.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.A.1. By König's infinity lemma [D25, 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.A.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 [JKK25] and joint work with Paul Knappe and Jan Kurkofka.

Every finite graph G can be decomposed in a canonical way that displays its local connectivity-structure [DJKK26]. These decompositions are defined via a suitable more tree-like covering of G, whose tangle-tree structure is projected down to G.

The covering graphs needed here are almost always infinite, and their tangle-tree structure is defined in terms of their (global) low-order separations. 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 establish the computability of the canonical graph-decompositions from [DJKK26].

2.1 Introduction

2.1.1 From tree-decompositions to graph-decompositions

Tree-decompositions are a well-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.

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 [C90]. In group theory, and in Bass-Serre theory more specifically, tree-decompositions are used to detect

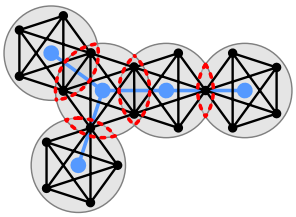

Figure 2.1: A tree-decomposition of a graph (black) into K_5's (grey).

product-structure in Cayley graphs such as amalgamated free products or HNN-extensions, for example in modern proofs of Stallings' Theorem [HLMR22; K10].

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 *tangle-tree-decomposition 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 tangle tree-decompositions. An example is depicted in Figure 2.1, where each K_5-subgraph corresponds to a tangle. The following recent work studies or constructs tangle tree-decompositions [A23; A24a; A25; CDHH14; CDHH16a; CDHH16b; C16; CDHS14; CG17; CK24; DEW19; DHL19; EKT21; EKT22b; E17; G16; JK24] or even more specific decompositions [BGHKMR23; CK23; EGL24; KP26]. Algorithmic applications include [GS15a; KEFKRTL23; K25], see [D24] for an overview.

Intuitively, tangle tree-decompositions 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 tangle tree-decompositions 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 [AKK25], see Figure 2.2c.

A recent development in graph minor theory, initiated by Diestel and Kühn [DK05] and followed up in [ADEFJKW23; AJKW25; DJKK26] and differently in [C22], is the study

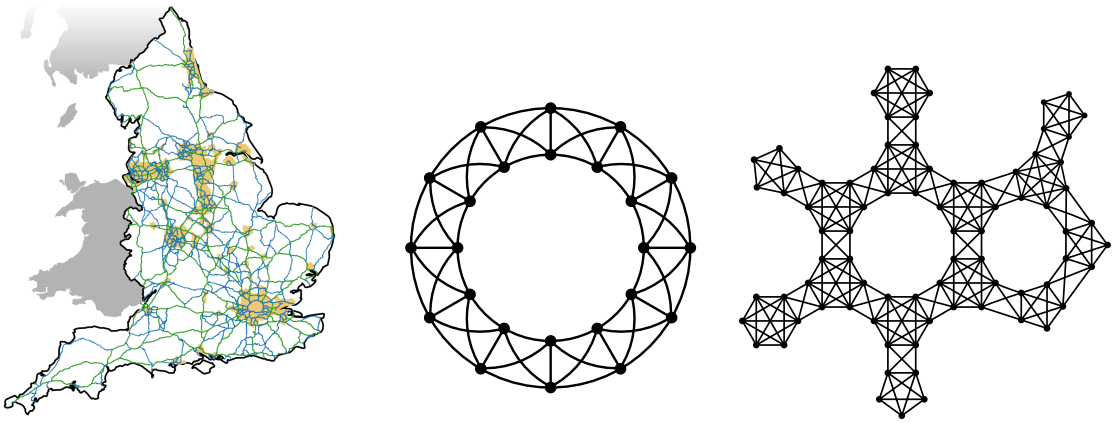

(a) The main road network of the UK

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

(c) A graph that resembles two of the authors' favourite molecule

Figure 2.2: Graphs whose connectivity-structure eludes tangle tree-decompositions

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 [DJKK26] have developed a 'graph-decomposition version' of the tangle tree-decomposition 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 $\mathcal{D}_r(G)$ of G in three steps.

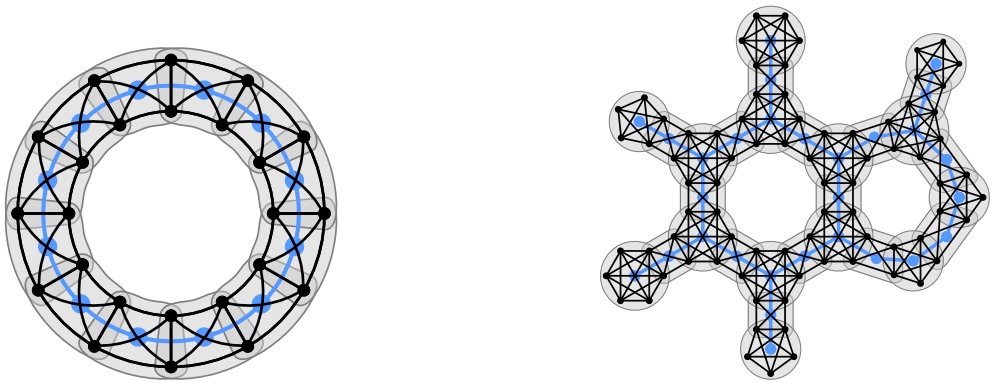

Figure 2.3: Graph-decompositions of black graphs into smaller parts (grey). The decomposition graphs are blue.

First, we take the *r-local covering* $p_r \colon G_r \to G$ of G introduced in [DJKK26], 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 tangle tree-decomposition of G_r. In the second step, we apply a tangle-tree-decomposition theorem for infinite graphs [CK24] to the covering G_r to obtain a tangle tree-decomposition 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 tangle tree-decomposition of G_r is 'folded back' via p_r to yield the graph-decomposition $\mathcal{D}_r(G)$ of G; see Construction 2.2.6. Following the above intuition, $\mathcal{D}_r(G)$ captures the arrangement of the r-locally highly connected regions in G.

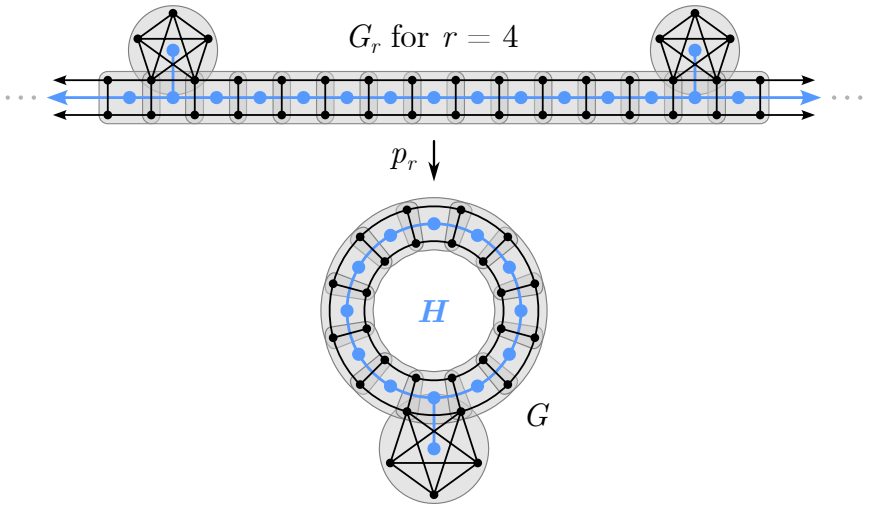

Figure 2.4: An H-decomposition of G (black) into 4-cycles and a K_5 (grey) over a cycle H (blue), obtained from a tree-decomposition of the r-local cover G_r for $r = 4$.

The graph-decomposition $\mathcal{D}_r(G)$ 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 $\mathcal{D}_r(G)$ 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 $\mathcal{D}_r(G)$ defaults to a tangle tree-decomposition of G.

The above construction of $\mathcal{D}_r(G)$ involves topology and infinite graphs. To compute such graph-decompositions, and thereby make them accessible to algorithms, we thus need to solve the following problem:

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

2.1.2 From separations to local separations

In a graph modelling a road network, small separators between large clusters form choke points that often get congested by traffic. But not all such choke points come from small separators: Consider for example the two red vertices in Figure 2.5a that block the two shortest routes between Birmingham (B) and Nottingham (N). Clearly, the two red vertices together form a choke point in between the two cities, as every detour is significantly longer. However, the two red vertices do not separate the entire road network, but only do so locally (which is enough to make them a choke point): they form an *r-local separator*, that is, a vertex-set that separates roughly a ball of given radius $r/2 > 0$ around it. In this chapter, we develop a formal way to study such choke points via local separators.

Carmesin [C22] initiated the study of r-local 1- and 2-separators, generalising the tree-like block-cutvertex decomposition and Tutte's 2-separator theorem [T84] (which SPQR-trees are bases on) 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 [C22] have since found applications to sparse networks and finite groups:

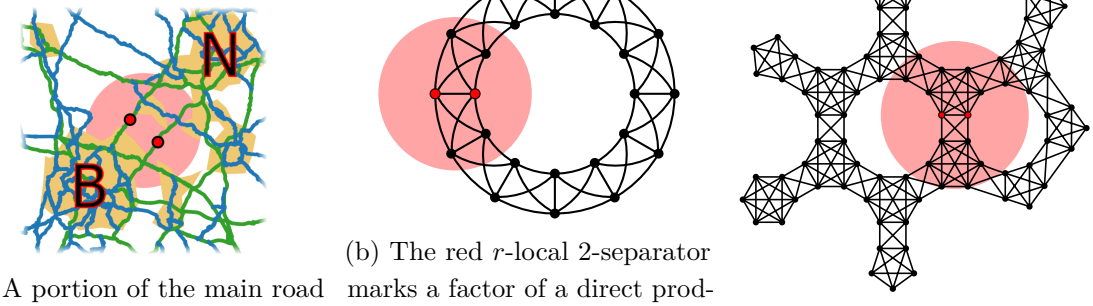

(a) A portion of the main road network of the UK.

(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.

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 [CKKT25] to prove a low-order Stallings-type theorem for finite nilpotent groups (Figure 2.5b).

Carmesin [C22, §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?*

2.1.3 Graph-decompositions and local separations

Our aim in this chapter is to present one unified solution to both Problem 1 and Problem 2: 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 $\mathcal{D}_r(G)$. Our work is fully compatible with both the existing perspectives of r-local coverings [DJKK26] and of r-local $(\leqslant 2)$-separations [C22].[1] This answers 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* $\Delta_r(G)$ to be the least distance between any two distinct vertices in G_r that p_r maps to the same vertex in G (setting $\Delta_r(G) := \infty$ if $G_r = G$).[2] 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 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. Perhaps surprisingly, the displacement is computable [AR]. We only use the

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

[2] Both G_r and p_r will be formally defined in Section 2.2. Equivalent definitions for $K(G, r)$ are discussed in Lemma 2.5.2.

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.

2.1.4 From trees of tangles to graphs of local bottlenecks

The work [EKT21; EKT22b] 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 tangle tree-decompositions.[3] As this construction works without referencing tangles, we refer to it as an *entanglement tree-decomposition*. The characteristic property of an entanglement tree-decomposition 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 $\mathcal{D}_r(G)$, which are obtained from tangle tree-decompositions 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 'r-local entanglement graph-decompositions'. Roughly speaking, an 'r-local entanglement graph-decomposition' 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 an entanglement tree-decomposition 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 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 to ensure that all minimum-order separations distinguishing any pair of tangles form a bottleneck.

As a main result, we show that every finite connected graph G has for every r and $k < K(G, r)$ an *r-local bottleneck graph-decomposition* $\mathcal{H}_r^{\leqslant k}(G)$. That is, $\mathcal{H}_r^{\leqslant k}(G)$ is a

[3] For ends instead of tangles, a similar construction but without entanglements had also been found by [R17].

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 [AKK25].

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 tangle tree-decomposition 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 tangle tree-decomposition $\mathcal{H}_r^{\leqslant k}(G_r)$ 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 $\mathcal{D}_r(G)$ 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 6. *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 canonical tangle tree-decomposition of scope 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 canonical 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 7. *There is an algorithm that determines $\mathcal{H}_r^{\leqslant k}(G)$ given G, r, k as in Theorem 6.*

Theorem 6 and Theorem 7 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.

ad (B): Our work has been used just recently by Köhl and Salarian [KS26] to establish combinatorial characterisations of virtually torsion-free and virtually free groups.

2.1.5 Organisation of the chapter

Section 2.2 recalls terminology about graph-decompositions. In Section 2.3, we introduce bottlenecks and generalise a canonical version of the tangle-tree-decomposition theorem from [EKT22b] 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 tangle-tree-decomposition 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 6 and 7. In Section 2.A, we discuss corners of local separations.

Acknowledgements

We are grateful to Nathan Bowler for pointing out an error in an early version. We thank Reinhard Diestel for sharing his very detailed feedback, which has improved the presentation across the board. We are especially grateful for his suggestion of a more elegant definition of bottlenecks. We thank Johannes Carmesin for his inspiration in the early stages of this project.

2.2 Review of graph-decompositions and coverings

Our notations and definitions follow [D17b]. 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 $(M_i : i \in I)$ of subsets M_i, possibly empty, that are disjoint and satisfy $M = \bigcup_{i\in I} M_i$. If $|I| = 2$, we speak of a *near-bipartition*.

We remark that most of what follows is a recap of definitions and concepts from [DJKK26] relevant to the present chapter; in [DJKK26] you will find proofs or reference to proofs.

2.2.1 Local coverings

We define a *walk* (of *length* k) in a graph G to be 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 $[W_1] \cdot [W_2] := [W_1 W_2]$ and $W_1 W_2$ 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 $W = W_0 Q W_0^-$. 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 C is a vertex set $\hat{X} \subseteq V(C)$ that contains for each $x \in X$ exactly one lift $\hat{x}$ of x. Equivalently, p restricts to a bijection $\hat{X} \to X$. A *lift* of a walk W in G is a walk in C which p projects to W.

There exists a covering $p_r\colon G_r \to G$ that 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 any lift of x_0. The covering $p_r\colon G_r \to G$ is unique up to isomorphism of coverings given G and r, and it is independent of the choice of x_0. In particular, the graph G_r is determined by just G and r. 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 $\Gamma_r(G)$. Since $\pi_1^r(G, x_0)$ is a normal subgroup of $\pi_1(G, x_0)$, this group $\Gamma_r(G)$ acts transitively on every fibre $p_r^{-1}(x)$ of the covering map p_r [H02, §1.3].

Lemma 2.2.1. [DJKK26, Lemma 4.6] *The binary cycle space of G_r is generated by the short cycles.*

We will often use this lemma implicitly when working with G_r.

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 ϱ is odd, then the $(\varrho/2)$-ball around v is precisely the subgraph induced on all vertices of distance at most $\varrho/2$ from v, while for even rho any edges between vertices of distance $\varrho/2$ from v will 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. Note that the surjectivity of these restrictions always holds and the definition requires additional injectivity.[4] By [DJKK26, Lemma 4.3 and 4.4], the r-local covering p_r may equivalently be described as the covering that is universal among the coverings $q\colon \hat{G} \to G$ preserving the $(r/2)$-balls. Thus, the r-local covering p_r maps every short cycle in G_r of length at most r isomorphically to G [DJKK26, Lemma 4.3]. Vice versa, any short cycle in G lifts to a short cycle in G_r.

[4] A notion that could be regarded as a precursor of the preservation of $(\varrho/2)$-balls has been studied by Georgakopoulos [G17].

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) separations, we denote the set of all orientations of 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[5]

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

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 then 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 automorphism group $\mathrm{Aut}(G)$ of G acts on the oriented separations of G by $(A, B) \mapsto (\varphi(A), \varphi(B))$ for $\varphi \in \mathrm{Aut}(G)$. This action is compatible with $\geqslant$. Thus, the respective action on the (unoriented) separations of G also preserves the notions of nested and crossing. We say that a set S of separations of G is *Γ-canonical* for some subgroup Γ of $\mathrm{Aut}(G)$ if S is invariant under the action of Γ on the separations of G. In the case of $\Gamma = \mathrm{Aut}(G)$, we just say that S is *canonical*.

The following definitions are illustrated in 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* of $\{A_1, A_2\}$ and $\{C_1, C_2\}$ is the intersection $X \cap Y$. The separator X is near-partitioned into its two X-links and the centre.

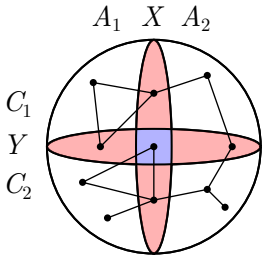

Figure 2.6: Four red links and a blue centre

[5]We remark that we use the partial order as defined in the 6th but not the 5th edition of Diestel's book [D17b]. That is, we write $\geqslant$ where readers familiar with previous chapters would expect $\leqslant$.

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 $(A_i \smallsetminus A_{3-i}) \cap (C_j \smallsetminus C_{3-j})$.*

(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 $E(X \cap Y, (A_i \smallsetminus A_{3-i}) \cap (C_j \smallsetminus C_{3-j}))$.*

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 $(A_i \smallsetminus A_{3-i}) \cap (C_j \smallsetminus C_{3-j})$. There is a vertex u of G outside of $(A_i \smallsetminus A_{3-i}) \cap (C_j \smallsetminus C_{3-j})$, 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 $E(X \cap Y, (A_i \smallsetminus A_{3-i}) \cap (C_j \smallsetminus C_{3-j}))$.

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

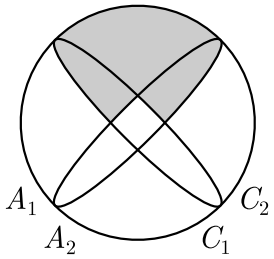

Figure 2.7: $(A_1, A_2) \geqslant (C_1, C_2)$ if and only if the grey area is empty. The grey area corresponds to $i = 1$ and $j = 2$ in Lemma 2.2.2 (ii)

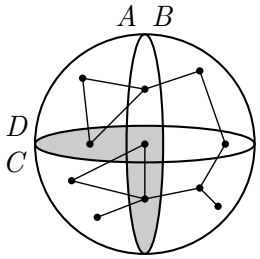

Figure 2.8: $\{B \cup D, A \cap C\}$ is one of the four corners of $\{A, B\}$ and $\{C, D\}$

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 $\{B \cup D, A \cap C\}$, $\{B \cup C, A \cap D\}$, $\{A \cup C, B \cap D\}$ and $\{A \cup D, B \cap C\}$. The corners $\{B \cup D, A \cap C\}$ and $\{A \cup C, B \cap D\}$ are *opposite*, and so are the corners $\{B \cup C, A \cap D\}$ and $\{A \cup D, B \cap C\}$. Any two corners that

are not opposite are *adjacent*. The two adjacent corners $\{B\cup D, A\cap C\}$ and $\{B\cup C, A\cap D\}$ are said to *lie on the same side* of $\{A,B\}$, namely *on the A-side*. Similarly, the other two adjacent corners $\{A\cup C, B\cap D\}$ and $\{A\cup D, B\cap C\}$ are said to *lie on the same side* of $\{A,B\}$, namely *on the B-side*.

The separations of G satisfy the following condition (F), see e.g. [D17b, Lemma 12.5.5].

(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*[6] on the separations of G if it satisfies

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

A set σ of oriented separations of G is a *star* if $(A,B)\geqslant(D,C)$ for every distinct two $(A,B),(C,D)\in\sigma$, and its *interior* $\operatorname{int}(\sigma)$ is $\bigcap_{(A,B)\in\sigma}B$. Let N be a nested set of proper 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 $(A,B)\in\sigma$. This $(A,B)\in\sigma$ is clearly unique.

Note that every $(A,B)\in\operatorname{ori}(N)$ lies in at most one splitting star. Moreover, if G is a finite graph, then the splitting stars of N form a partition of $\operatorname{ori}(N)$ [D18b, §3]. Kneip and Gollin [GK21, Corollary 3.3] showed that this also holds for infinite graphs, if additionally $\operatorname{ori}(N)$ has no $(\omega+1)$-chain.

2.2.3 Decompositions

In this section, we mostly follow [DJKK26]. Let G be a graph, and H a multigraph. Let $\mathcal{V}=(V_h : h\in V(H))$ be a family of sets V_h of vertices of G indexed by the nodes of H. The pair $(H,\mathcal{V})$ is an *H-decomposition* or *graph-decomposition* of G if

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

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

We refer to the V_h as its *bags* and to the W_v as its *co-bags*, and we call H its *decomposition graph*. If H is a tree, then a graph-decomposition $(H,\mathcal{V})$ is a *tree-decomposition*.

Note the duality $v\in V_h \iff h\in W_v$. We may thus reconstruct the families $\mathcal{V}=(V_h : h\in V(H))$ and $\mathcal{W}=(W_v : v\in V(G))$ from 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

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

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

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

These G_h, then, are our chosen *parts* of the decomposition $(H, \mathcal{V})$, and 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'). We often introduce graph-decompositions (with parts) as a pair $(H, \mathcal{V})$ (or $(H, \mathcal{G})$) and tacitly assume that $\mathcal{V} = (V_h : h \in V(H))$ (or $\mathcal{G} = (G_h : h \in V(H))$), respectively.

In the case of tree-decompositions $(T, \mathcal{V})$, we generally allow parts G_t other than the induced subgraphs $G[V_t]$ on the bags. However, throughout this chapter, parts of a tree-decomposition will always be the induced subgraphs $G[V_t]$. This agrees with the traditional definition of parts of tree-decompositions, as for example in [D17b].

For a category-theoretic perspective on graph-decompositions and tree-decompositions, see [BFT25; BKMM25].

Construction 2.2.3 (Tree-decompositions from nested sets of 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. T is a graph whose nodes are the splitting stars of N and whose edges 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 bag V_t of a node t of T is $\mathrm{int}(t)$, with t viewed as a splitting star of N, and we usually flesh out its part as $G_t := G[V_t]$.

If N is finite, then T is a tree; moreover, $(T, \mathcal{V})$ is a tree-decomposition of G [D18b, §6]. If N is infinite, elements of $\mathrm{ori}(N)$ may fail to lie in a splitting star, and thus T may not be a tree. But even if T is a tree, then the candidate $\mathcal{T}$ need in general not be a tree-decomposition of a connected and locally finite graph G [EKT22b, Example 4.9]. However, we have

Lemma 2.2.4. [DJKK26, Lemmas 7.2 & 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. Moreover, if N is Γ-canonical for some subgroup Γ of $\mathrm{Aut}(G)$, then so is $\mathcal{T}(N)$.*

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 $i \in \{1, 2\}$ 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 $A_i^e := \bigcup_{t \in V(T_i)} V_t$ and T_i is the component of $T - t_1 t_2$ which contains t_i, see [D17b, 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 $(\mathrm{ori}(N), \geqslant)$. The splitting stars of N correspond, under this isomorphism, to the stars of all edges of T at a fixed vertex, oriented towards that vertex.*

Proof. This is well known; see for example [EKT22b, Proof of Lemma 2.7]. □

An *action* of a subgroup Γ of $\mathrm{Aut}(G)$ on a graph-decomposition $(H, \mathcal{V})$ of G is an action of Γ on the graph H as automorphisms of H that commutes with $h \mapsto V_h$, i.e. every $\varphi \in \Gamma$ acts on H as an automorphisms of H so that $\varphi(V_h) = V_{\varphi(h)}$ for all $h \in V(H)$. If such an action exists, then $(H, \mathcal{V})$ is Γ-*canonical*. In the case of $\Gamma = \mathrm{Aut}(G)$, we just say that $(H, \mathcal{V})$ is *canonical*. If the graph-decomposition $(H, \mathcal{V})$ comes with parts G_h, then it is Γ-*canonical* if the action also respects these parts in that $\varphi(G_h) = G_{\varphi(h)}$ for all $\varphi \in \Gamma$ and $h \in V(H)$, and in the case of $\Gamma = \mathrm{Aut}(G)$, we call it *canonical*.

If $\mathcal{T}(N)$ is a tree-decomposition and N is Γ-canonical, then the $\varphi \in \Gamma$ act on $\mathrm{ori}(N)$ as bijections $(A, B) \mapsto (\varphi(A), \varphi(B))$, so the action is compatible with $\geqslant$. Thus, $\mathcal{T}(N)$ is Γ-canonical by construction (which immediately yields the 'moreover'-part of Lemma 2.2.4). Moreover, every automorphism φ of G commutes with at most one automorphism ψ of the decomposition tree T in that $\varphi(V_t) = V_{\psi(t)}$ [DJKK26, 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 as automorphisms of T that commutes with $t \mapsto V_t$.

Construction 2.2.6 (Graph-decompositions via local coverings [DJKK26, Construction 3.8])**.**
Let G be a connected graph and let $r \in \mathbb{N}$. Let $\Gamma_r := \Gamma_r(G)$ be the group of deck transformations of the r-local covering of G, and let $\mathcal{T} = (T, \mathcal{V})$ be a Γ_r-canonical tree-decomposition of G_r. We define the graph-decomposition $\mathcal{H} = (H, \mathcal{V})$ of G, 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 bag V_h of a node h of H is defined as $V_h := p_r(V_t)$ for $t \in h$. If $(T, \mathcal{V})$ comes with parts G_t, then we flesh out the family $\mathcal{G}$ of parts G_h of $(H, \mathcal{V})$ as $G_h := p_r(G_t)$ for $t \in h$. We say that $(H, \mathcal{V})$ and $(H, \mathcal{G})$ are *defined* by $(T, \mathcal{V})$ (and its parts) *via* p_r, or that $(H, \mathcal{V})$ and $(H, \mathcal{G})$ are *obtained* from $(T, \mathcal{V})$ (and its parts) by *folding via* p_r.

Pairs $(H, \mathcal{V})$ from Construction 2.2.6 are graph-decompositions by [DJKK26, 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; EKT22b] 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 tangle tree-decompositions, see Construction 2.3.6. This in turn makes possible the formulation of an r-local version of a tree of tangles that is compatible with the guarantee $K(G, r)$, see Section 2.7. We do not develop a local version of tangles, as the need did not arise in this project. Still, we think it would be exciting to have a local version of tangles.

2.3.1 Bottlenecks

For $n \in \mathbb{N}$, we denote $\{1, \dots, n\}$ by $[n]$. A *$\top$-star* of G 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 $B_i = A_j \cup A_k$ whenever $\{i, j, k\} = [3]$. 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* $(X_i \cap X_j) \setminus Z = X \setminus X_k$ by X_{ij} whenever $\{i, j, k\} = [3]$.

These definitions are inspired by their usual occurrence in a corner diagram:

Example 2.3.1. This example is depicted in 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. The 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.

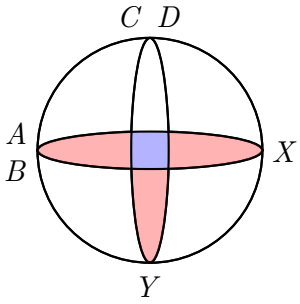

Figure 2.9: The situation in Example 2.3.1: the three ij-links are red, the centre of σ is blue

It is immediate from the definition of $\top$-star that

$$X_i = X_{ij} \sqcup Z \sqcup X_{ik} \text{ whenever } \{i, j, k\} = [3] \tag{2.1}$$

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 $G[A_i \setminus B_i]$ with $x, y \in N(K)$; see Figure 2.10.

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.*

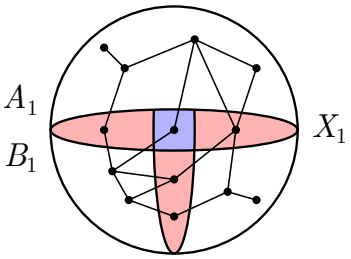

Figure 2.10: (A_1, B_1) is part of a $\top$-star that is relevant with base $\{A_1, B_1\}$. The links are red and the centre is blue. The red link on the right is the 23-link

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 their two oriented 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 $Z := \bigcap_{i \in [3]} X_i$ satisfies $Z = A \cap B \cap C \cap D$ by the definition of σ. Hence, $N(K) = X_{23} \cup Z \cup Y$, where Y is the link $(C \cap D) \setminus B$. 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 $G[C \setminus D]$ and has a neighbour $x \in X_{23}$, it must also be included in $G[A_2 \setminus B_2]$. Thus, K is a component of $G[A_2 \setminus B_2]$ with a neighbour in Z.

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

Finally, we assume that X_{12} is non-empty while Y is empty. Lemma 2.3.2 (applied to $(A, B) := (D, C)$, $X := A \cap B$, $(C, D) := (A, B)$ and $Y := C \cap D$) yields that the centre Z is non-empty. Consider an x–Z path Q through the component $K \subseteq G[C \setminus D]$. 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 satisfying:

(B) Whenever β contains the 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. For every pair of distinguishable tangles τ, τ' in a graph G, the set $\beta = \beta(\tau, \tau')$ of all separations efficiently distinguishing τ and τ' forms a bottleneck in G; see [D17b, §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 the same order k. We claim that β is k-bottleneck. By their efficiency, every separation in β is tight; compare [EKT22b, Lemma 6.1 and its proof]. 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 for τ', 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 the tangle property for τ. 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, $\{A_j, B_j\} \in \beta$. □

An oriented separation (A, B) *points towards* a vertex set Y of G if $Y \subseteq B$. Let Y be a clique in G of size k, i.e. a set of vertices which induces a complete subgraph in G. Then precisely one orientation of a given separation $\{A, B\}$ of order $< k$ points towards Y. By τ_Y, we denote the set of oriented separations of order $< |Y| = k$ pointing towards Y. We warn the reader that τ_Y might not be a tangle of order k. A separation $\{A, B\}$ of G *distinguishes* two cliques Y, Z of G if $(A, B) \in \tau_Y$ and $(B, A) \in \tau_Z$. A separation distinguishes Y and Z *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 and consisits of 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, let $\{A, B\}$ be a separation in β. We may

assume that $(A, B) \in \tau_Y$ and $(B, A) \in \tau_Z$. Then each vertex in $Y \cap (A \setminus B)$ from each vertex in $Z \cap (B \setminus A)$ order-minimally. Since $k < \min\{|Y|, |Z|\}$ by assumption, both these sets are non-empty. 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 $Y \subseteq A_1$ and $Z \subseteq B_1$; in particular, $Y \subseteq B_2, B_3$. Since the order of σ is $< \min\{|Y|, |Z|\}$, the clique Z does not lie in the union of the separators of the (A_i, B_i). Thus, $\{A_1, B_1\} \in \beta$ yields that either A_2 or A_3 contains Z, and thus $\{A_2, B_2\}$ or $\{A_3, B_3\}$, respectively, is in β. □

2.3.2 A tree of bottlenecks

The following construction transfers the construction introduced in the proof of [EKT22b, Theorem 1.2] to bottlenecks:

Construction 2.3.6 (Nested separations from bottlenecks)**.** Let G be a graph, and let $\mathcal{B}$ be a set of bottlenecks in G. For every $k \in \mathbb{N}$, let $\mathcal{B}^k(G)$ be the set of all k-bottlenecks in $\mathcal{B}$ in G and $S^k := \bigcup \mathcal{B}^k(G)$. Recursively for all $k \in \mathbb{N}$, define nested sets $N^k(G) \subseteq S^k$ of separations, as follows. Let $x^k \colon S^k \to \mathbb{N}$ assign to each separation $s \in S^k$ the number of separations in S^k that s crosses, assuming for now that this number is finite. 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 $N(G) := \bigcup_{k \in \mathbb{N}} N^k(G)$.

Whenever we use the notation of $N^k(G)$ or $N(G)$ without reference to any set $\mathcal{B}$ of bottlenecks, we set $\mathcal{B}$ to be the set of all bottlenecks in G.

Theorem 2.3.7. *Let G be a locally finite, connected graph. Given any set $\mathcal{B}$ of bottlenecks in G, 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. Moreover, if $\mathcal{B}$ is Γ-invariant for a subgroup Γ of $\mathrm{Aut}(G)$, then $N(G)$ and $N^{<k}(G)$ (for $k \in \mathbb{N}$) are Γ-canonical.*

Here, a set $\mathcal{B}$ of bottlenecks is Γ-*invariant* for a subgroup Γ of $\mathrm{Aut}(G)$ if the action of Γ on the bottlenecks of G (which is inherited from the action of $\mathrm{Aut}(G)$ on the separations of G) is invariant on $\mathcal{B}$.

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 [H91a, p. 2.4] as in the proof of [EKT22b, 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 adjacent corners c, d of order $\leqslant k_2$. 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, as shown above, c has order $< k_1$. □

Lemma 2.3.10. *For every k-bottleneck β in any given set $\mathcal{B}$ of bottlenecks in a locally finite, connected graph G, the set $N^k(G, \beta)$ is non-empty.*

Proof. It suffices to show for $k > 0$ that β contains separations that are nested with $N^{<k}(G)$. 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 (F1) and the fact that $N^{<k}(G)$ is nested. □

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), in the same way as in the proof of [CK24, Corollary 3.3] or [D17b, Lemma 12.5.6]. □

Proof of Theorem 2.3.7. Lemma 2.3.8 ensures that x^k in Construction 2.3.6 assigns finite values. 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 $\leqslant k$. By symmetry, we may assume without loss of generality that the corners c, c', d' depicted in Figure 2.11 have order $\leqslant k$.

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 $N^k(G, \beta)$.

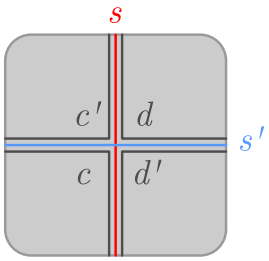

Figure 2.11: The four corners in the proof of Theorem 2.3.7

Case 2: The corner d has order $\leqslant k$. Applying Lemma 2.3.3 and (B) to both sides of both $s \in \beta$ and $s' \in \beta'$ yields that β meets $\{c, c'\}$ and $\{d, d'\}$ and β' meets $\{c, d'\}$ and $\{c', d\}$. If for a pair of opposite corners one is in β and the other in β', then we obtain a contradiction from (2.2) as in Case 1. So we may assume that this does not happen. By symmetry, we may assume without loss of generality that $c \in \beta$, and hence $d \notin \beta$. Since β' meets $\{c', d\}$, we have $c' \in \beta'$, which in turn gives $d \in \beta$, and finally $d' \in \beta'$. By symmetry, we may assume without loss of generality $x^k(s) \leqslant x^k(s')$. 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 $N^k(G, \beta')$.

The 'moreover'-part is immediate from the canonicity of Construction 2.3.6. □

A *tree of tangles* N of G of *scope* k is a nested set of separations of G which contains, for every pair of $(\leqslant k)$-distinguishable tangles τ_1, τ_2 in G, a separation that efficiently distinguishes τ_1 and τ_2. A *tangle tree-decomposition* $\mathcal{T}$ of G of *scope* k is a tree-decomposition of G such that the set of induced separations of $\mathcal{T}$ is a tree of tangles of G of scope k.

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

Proof. Apply Theorem 2.3.7 with $\mathcal{B}$ the set of all bottlenecks in G. Then Example 2.3.4 yields that $N^{\leqslant k}(G)$ is a canonical tree of tangles. By Lemma 2.2.4, $\mathcal{T}(N^{\leqslant k}(G))$ is a canonical tree-decomposition. □

A tree of tangles that has scope k for every $k \in \mathbb{N}$ has *infinite* scope. Likewise, a tangle tree-decomposition that has scope k for every $k \in \mathbb{N}$ has *infinite* scope.

Theorem 2.3.13. *Let G be finite connected graph. Then $N(G_r)$ is a canonical tree of tangles of the r-local covering G_r of infinite scope. Moreover, $\mathcal{T}(N(G_r))$ is a canonical tangle tree-decomposition of G_r of infinite scope.*

Proof. Again apply Theorem 2.3.7 with $\mathcal{B}$ the set of all bottlenecks in G. Then Example 2.3.4 yields that $N(G_r)$ is a canonical tree of tangles. To see that $\mathcal{T}(N(G_r))$ is indeed a canonical tree-decomposition, we combine Lemma 2.2.4 with the following results of [DJKK26]: Theorem 3, Theorem 6.2, Lemma 6.3 and Lemma 8.1. □

2.4 Local separations

In this section, we define *r-local separations* formally, along with notions of *r-local separators* and *r-local components*. Ideally, we would like the r-local notions to be compatible with the perspective of r-local coverings in two ways: On the one hand, each r-local notion should generalise its global counterpart so that both notions are equivalent in all r-local covers. On the other hand, each r-local notion should get along with every r-local covering $p_r\colon G_r \to G$ in that we can move in between the r-local separations (say) of G and the r-local separations of G_r by means of lifting or projecting.

Our notions will not satisfy these properties in general, but we shall find conditions under which they do and which cover all relevant cases. The equivalences between the r-local notions and their global counterparts in r-local covers are shown in Correspondences 2.4.1, 2.4.4 and 2.4.6. Projection 2.4.13 and Lift 2.4.16 show how r-local components can be lifted and projected, respectively, between graphs their r-local covers. Projection 2.4.18 and Lift 2.4.19 achieve this for r-local separations, which includes r-local separators as a special case.

2.4.1 Local separations and their counterparts in the local cover

Let G be a connected graph and $\emptyset \neq X \subseteq V(G)$. 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. As G is connected, the components C of $G - X$ partition ∂X into the componental cuts $E_G(C, X)$. 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 (see Section 2.2.1) in G which starts with e and ends with f. The notion of *r-local components* below is motivated by this equivalence relation.

An *X-edge* is an edge with both endvertices in X; note that X-edges are not in ∂X. An *r-local X-walk* in G is an X-walk in some short cycle in G, one of length $\leqslant r$, 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

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

The r-local components at X thus further partition the componental cuts at X. Every r-local X-walk meets every r-local component at X in precisely its first and last edge or in none, and thus

$$\text{every short cycle } O \text{ meets every } r\text{-local component at } X \text{ in an even number of edges.} \tag{2.4}$$

We remark that, if a short cycle O meets X in two vertices, then O splits into edge-disjoint r-local X-paths.

We call $X \subseteq V(G)$ an *r-local separator* of G if there are at least two r-local components at X. For the right intuition, note that the qualification of being 'r-local' refers to the way in which X 'separates' (namely, only locally), not to the position of X in G. In particular, every separator of G is also an r-local separator. Figure 2.12 shows an example of an r-local separator that is not a (global) separator. The size of an r-local separator, which we note cannot be zero, is its *order*. An r-local separator of order k is also called an *r-local k-separator*.

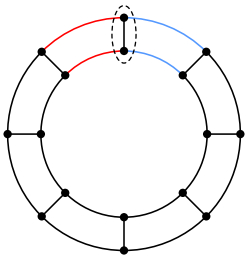

Figure 2.12: There are two 4-local components (red and blue) at X (encircled)

Our aim in introducing local separations was that, for every r, the relevant r-local k-separations of G should lift to (global) k-separations of G_r. Before we discuss lifting properties, let us show that the r-local separations of G_r (to which we hope our r-local separation of G will lift) are indeed (global) separations of G_r – as long as they are relevant in the following natural sense of being 'locally tight'.

Let us call an r-local component F at X *tight* if every vertex in X is incident with some edge in F, and let us call an r-local separator *tight* if there are at least two tight r-local components at X. Note that although every separator is also an r-local separator, it can be tight as a separator but not as an r-local separator, see Figure 2.13. Conversely, an r-local separator need not be a separator, but if it is, then the tightness of the r-local separator clearly implies the tightness of the separator. In the r-local cover these notions coincide:

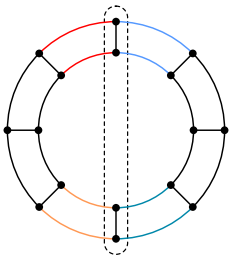

Figure 2.13: The encircled vertex set X is a tight separator of G. However, there are four 4-local components at X, none of which is tight. So while X is a 4-local separator, it is not a tight one.

Correspondence 2.4.1 (Tight separator)**.** *Assume that the binary cycle space of a connected graph 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 $\emptyset \neq X \subseteq V(G)$ we denote by $\mathcal{W}_r(X)$ the set of all walks that are concatenations of r-local X-paths. 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.

Lemma 2.4.2. *Every tight r-local separator of a connected graph G is r-tomic.*

Proof. This is immediate from the definitions. □

Lemma 2.4.3. *Assume that the binary cycle space of a connected graph 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 2.4.4 (Local components)**.** *Assume that the binary cycle space of a connected graph G is generated by cycles of length $\leqslant r$. The r-local components at an r-tomic set $X \subseteq V(G)$ are precisely the 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 always 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 W which passes through E oddly often. Consider the element W' of the binary cycle space induced by the closed walk W: an edge of G maps to 0 if W passes it evenly often, and otherwise, the edge maps to 1; in particular, the values of W' on the edges in E sum to an odd number. Since the binary cycle space of G is generated by short cycles, W' is the symmetric difference of the edges sets of short cycles $O_1, \dots, O_n$. Hence, some short cycle O_i meets E oddly, and therefore contradicts (2.4). □

Proof of Correspondence 2.4.1. 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 2.4.3 or Lemma 2.4.2, respectively. Hence, Correspondence 2.4.4 completes the proof. □

Lemma 2.4.5. *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 2.4.3, 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 $x \in Y' := X \setminus Y = X_{23} \setminus Y$. 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 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 these 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 an X_{23}–X_{12} edge, then this X-edge witnesses that some vertex in Y' is in a common r-tom with some vertex in $X_{12} \subseteq Y$, which contradicts the maximality of Y as 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), 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 $E(X_2 \cap Y, A_2 \setminus B_2)$. 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 maximality of Y as an r-tom of X. □

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. We use this perspective to introduce local separations.

An *r-local separation* of G is a triple $\{E_1, X, E_2\}$ where X is a non-empty set of vertices of G and E_1, E_2 near-bipartition of ∂X such 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.

An oriented separation (A_1, A_2) of G *induces* the oriented r-local separation (E_1, X, E_2) with $X := A_1 \cap A_2$ and $E_i := E_G(X, A_i \setminus X)$. 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. If an r-local separation is introduced as $\{E_1, X, E_2\} := s_r$ when $\{A_1, A_2\} = s$ was previously given, we tacitly assume that $E_i = E_G(X, A_i \setminus X)$ and not $E_i = E_G(X, A_{3-i} \setminus X)$. We remark that by definition $s \mapsto s_r$ is an injective map from the (oriented) separations of G to the (oriented) r-local separations of G.

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 2.4.4. □

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 encircled vertices, 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.

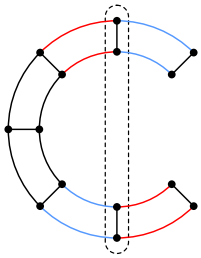

Figure 2.14: An r-local separation that is not induced by a separation, as the sides of the r-local separation restricted at the r-local atoms of its separator are flipped.

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 2.4.3 or Lemma 2.4.2, so the other of s, s_r is tight as well by Correspondence 2.4.4. 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

Let X be a vertex set of G. If X is r-tomic, then the deck transformations of p_r act transitively on the set of r-toms of $p_r^{-1}(X)$ in G_r by $Y \mapsto \gamma(Y)$ for any deck transformation γ. In Figure 2.15, the set X of two vertices of G is 4-tomic. It's preimage $p_4^{-1}(X)$ has infinitely many 4-toms in the 4-local cover G_4 of G.

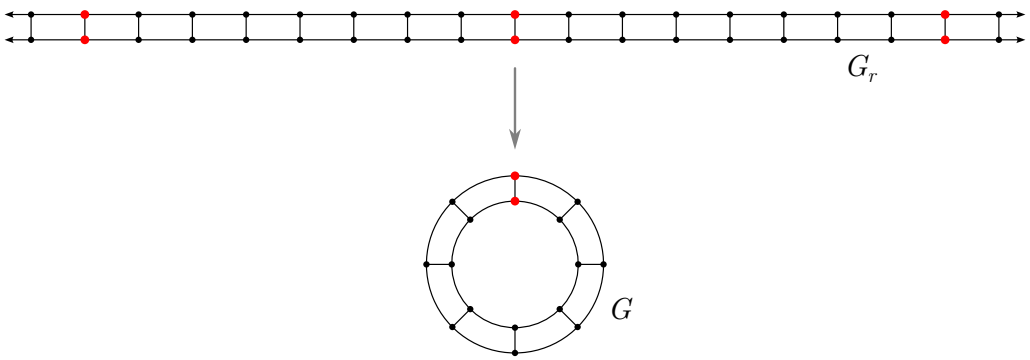

Figure 2.15: The set X is red in G (bottom) and $p_4^{-1}(X)$ is red in G_4 (top).

Projection 2.4.9 (Locally atomic)**.** *Let $\hat{X}$ be a vertex set of G_r. If $\hat{X}$ is r-tomic, then its projection $p_r(\hat{X})$ is r-tomic.*

Proof. Let $\hat{x}$ and $\hat{y}$ be two arbitrary vertices in $\hat{X}$ that are linked by an r-local $\hat{X}$-path $\hat{W}$ in G_r. Since p_r maps short cycles of G_r isomorphically to short cycles in G, the projection $W := p_r(\hat{W})$ is a concatenation of r-local X-paths in G that joins $x := p_r(\hat{x})$ to $y := p_r(\hat{y})$. As we have chosen $\hat{x}, \hat{y}$ arbitrarily in $\hat{X}$, it follows immediately that $p_r(\hat{X})$ is r-tomic. □

Projection 2.4.10 (Local atoms)**.** *Let X be a vertex set of G. Then p_r maps the r-toms of $\hat{X} := p_r^{-1}(X)$ to the r-toms of X. In particular, if X is r-tomic, then p_r maps the r-toms of $\hat{X}$ onto X.*

Proof. Let $\hat{Y}$ be an r-tom of $\hat{X}$. By Projection 2.4.9, its projection $Y := p_r(\hat{Y})$ is r-tomic. Suppose Y is not an r-tom of X, i.e. there is an r-local X-path W in G that joins a vertex $y \in Y$ to a vertex $x \in X \setminus Y$. Since any short cycle of G lfits to a short cycle of the r-local cover G_r, the r-local X-path W lifts to an r-local $\hat{X}$-path $\hat{W}$, as $\hat{X} = p_r^{-1}(X)$. But $\hat{W}$ witnesses that $\hat{Y}$ is not an r-tom of $\hat{X}$, which is a contradiction. □

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 closed* (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 closed, then p_r maps edges of $\partial\hat{X}$ to ∂X (rather than to X-edges).

Observation 2.4.11. *For every $X \subseteq V(G)$, the r-local atoms of $p_r^{-1}(X)$ are its $\subseteq$-minimal non-empty r-locally closed subsets. More generally, $\hat{X} \subseteq V(G_r)$ is r-locally closed if and only if it is a union of r-local atoms of $p_r^{-1}(p_r(\hat{X}))$.*

Lift 2.4.12 ($e \sim_r f$ at X in G)**.** *Given a connected graph G, let $\hat{X}$ be an r-locally closed 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 $\hat{f} := \hat{e}$. Otherwise, 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 $\hat{e}$(which is a short cycle in G_r), 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 is the last edge of $\hat{W}$. As $\hat{X}$ is r-locally closed, $\hat{W}$ cannot end in $p_r^{-1}(X) \setminus \hat{X}$: it 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). Since p_r maps short cycles of G_r isomorphically to short cycles of G, 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 closed. □

Projection 2.4.13 (Local components)**.** *Given a connected graph G, let $\hat{X}$ be an r-locally closed 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 closed, all of C is included in ∂X. Then C is included in an r-local component C' at X by Lift 2.4.12 (iii)→(i). Conversely, C' is included in C by Lift 2.4.12 (i)→(ii), so $C = C'$ as desired. □

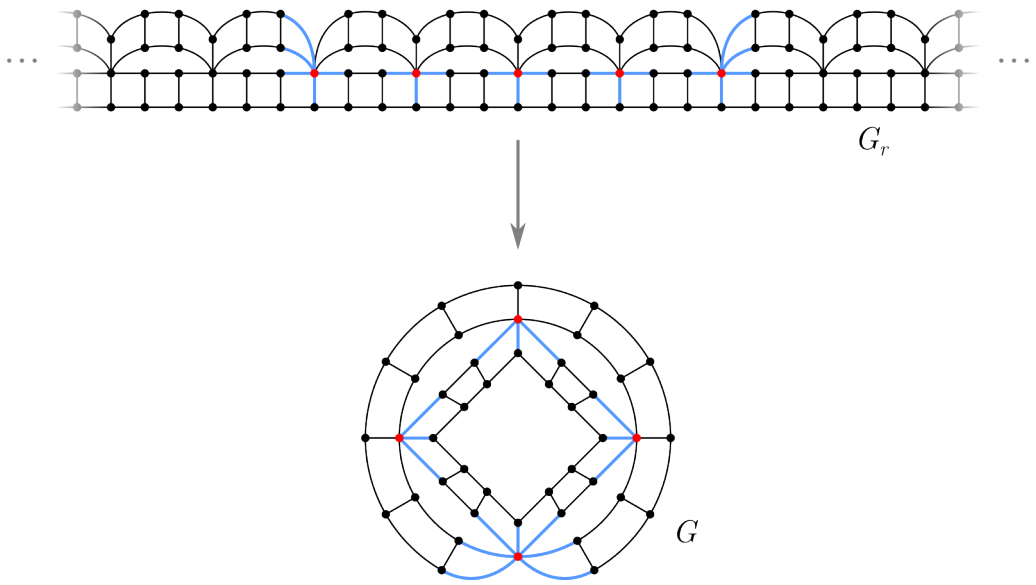

Figure 2.16: Let $r := 6$. The figure shows an r-local component (blue, top) at a non-r-locally-clopen vertex set $\hat{X}$ (red) in G_r whose projection to G (blue, bottom) fails to be an r-local component at $p_r(\hat{X})$ (red)

Example 2.4.14. The assumption in Projection 2.4.13 that $\hat{X}$ is r-locally closed is necessary, see Figure 2.16.

Projection 2.4.13 implies that for every r-local separation (E_1, X, E_2) of G there is an r-local separation $(\hat{E}_1, \hat{X}, \hat{E}_2)$ of G_r such that $p_r(\hat{X}) = X$ and $p_r(\hat{E}_i) = E_i$ for $i = 1, 2$: just take $\hat{X} := p_r^{-1}(X)$ and $\hat{E}_i := p_r^{-1}(E_i)$ for $i = 1, 2$.

If X is r-tomic, we can find such an $(\hat{E}_1, \hat{X}, \hat{E}_2)$ with $\hat{X}$ r-tomic too: just take as $\hat{X}$ an r-tom of $p_r^{-1}(X)$ instead of the entire set $p_r^{-1}(X)$ (cf. Projection 2.4.10), and define $\hat{E}_i := p_r^{-1}(E_i) \cap \partial\hat{X}$ for $i = 1, 2$, i.e. an r-local component $\hat{C}$ at $\hat{X}$ is in $\hat{E}_i$ if its projection $p_r(\hat{C})$ is in E_i. We have thus shown the following:

Corollary 2.4.15. *Every r-local separation (E_1, X, E_2) of G is the projection of some r-local separation $(\hat{E}_1, \hat{X}, \hat{E}_2)$ of G_r. If X is r-tomic, then $\hat{X}$ can be chosen as any r-tom of $p_r^{-1}(X)$.* □

Even if the set X in Corollary 2.4.15 is r-tomic, the r-tomic set $\hat{X}$ can be much larger than X. Figure 2.17 below shows an example where $\hat{X}$ is infinite but X is finite. Also G_r has more r-local components at $\hat{X}$ than G has at X in this example. Hence, G_r has r-local separations with separator $\hat{X}$ that do not project to r-local separations of G.

However, it turns out that the latter problem will go away when we solve the former: if we require that $\hat{X}$ is no larger than X, then r-local separations of G_r with separator $\hat{X}$ do project to r-local separations of G (Projection 2.4.18).

Recall that a lift of $X \subseteq V(G)$ to G_r is 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 closed. However, not every $X \subseteq V(G)$ has a lift that is r-locally closed. Sufficient conditions will be discussed in Section 2.5.

Lift 2.4.16 (Local components)**.** *Let X be a set of vertices in a connected graph G, and let $\hat{X}$ be an r-locally closed 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 closed, Projection 2.4.13 yields that $C := p_r(\hat{C})$ is an r-local component at X in G. The assumption that $\hat{X}$ is a lift of X implies by the definition of coverings that the restriction of p_r to $\hat{C}$ is also injective, giving (i), and that these bijections $\hat{C} \to C$ satisfy (ii).

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

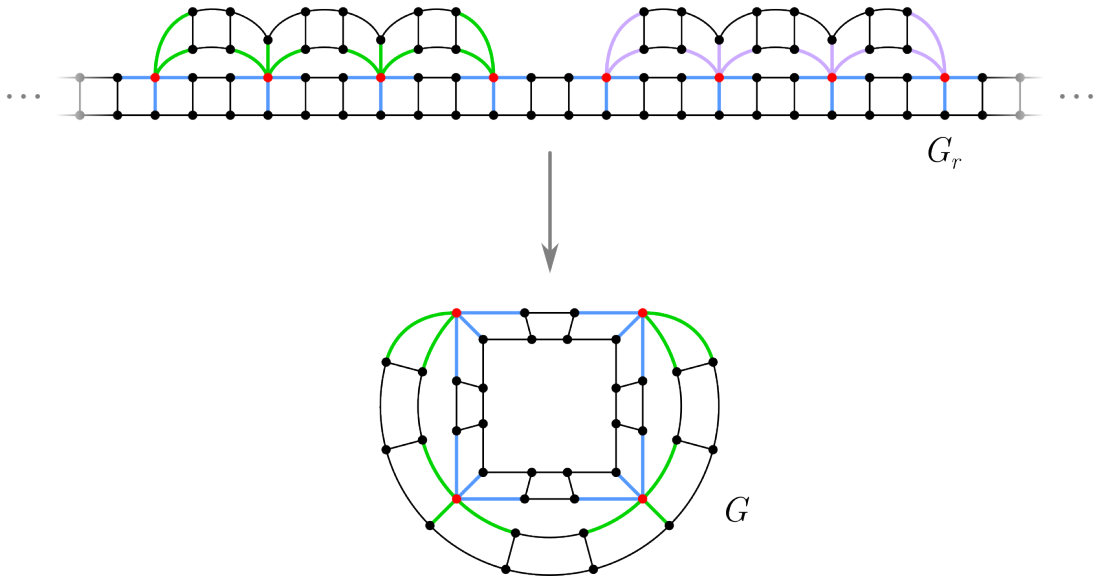

Figure 2.17: Let $r := 8$. The figure shows an r-locally closed vertex set $\hat{X}$ (red, top) in G_r which is not a lift of $X = p_r(\hat{X})$ (red, bottom). The coloured edge sets in G_r are r-local components at $\hat{X}$, and the coloured edge sets in G are r-local components at X.

Example 2.4.17. The assumption in Lift 2.4.16 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 2.17. Let X be the red set of vertices in G. Then $\hat{X} := p_r^{-1}(X)$ is the only non-empty r-locally closed 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 2.4.16 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 2.4.18 (Local separations)**.** *Let X be a set of vertices in a connected graph G, and let $\hat{X}$ be an r-locally closed lift of X to G_r. Suppose $\hat{X}$ is the separator of an r-local k-separation $\{\hat{E}_1, \hat{X}, \hat{E}_2\}$ 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 2.4.16. □

In the context of Projection 2.4.18, when $\{\hat{E}_1, \hat{X}, \hat{E}_2\}$ is introduced as $\hat{s}$, we often denote $\{E_1, X, E_2\}$ by $p_r(\hat{s})$ and refer to it as the *projection* of $\hat{s}$, and similarly for oriented local separations $(\hat{E}_1, \hat{X}, \hat{E}_2)$.

Lift 2.4.19 (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 closed 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 2.4.16. □

In the context of Lift 2.4.19, 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) is an $\hat{X}$-lift of (E_1, X, E_2) for some $\hat{X}$ as in Lift 2.4.19. We emphasise that if $\hat{X}$ is not only r-locally closed but even r-tomic, then, by Correspondence 2.4.6, a lift $\hat{s}$ of an r-local separation s of G is induced by a (global) separation of G_r. Usually, we will later find $\hat{X}$ as an r-tom of $p_r^{-1}(X)$ via Lemma 2.5.7 *.

Projection 2.4.18 and Lift 2.4.19 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 closed 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 Displacement: a guarantee for the existence of lifts

While Projection 2.4.18 and Lift 2.4.19 explain the interplay of r-local separations with the covering map p_r, they both rest on the assumption that the (r-local) separator $\hat{X}$ of G_r is an r-locally closed lift of the r-local separator X of G. But not all r-local separators X admit such r-locally closed lifts $\hat{X}$. For example, it is easy to check that every lift of the r-local 4-separator of G shown in Figure 2.16 fails to be r-locally closed in G_r.

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 $|X| < 2\Delta_r(G)/r$,

then X has an r-locally closed lift $\hat{X}$ (Lemma 2.5.7*). This then yields the desired lifting/projection lemmas for r-local separations, Projection 2.5.9* and Lift 2.5.10*. We also give various characterisations of the displacement (Lemma 2.5.2) as well as a simple combinatorial lower bound (Lemma 2.5.3). Finally, we provide a construction scheme for graphs (Examples 2.5.11 and 2.5.13) that 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

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 [DJKK26, Lemma 4.3] shows this. □

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 [DJKK26, Lemma 4.2].

(ii)→(iii). Since the covering p_r preserves $(\varrho/2)$-balls, the cycles in G of length $\leqslant \varrho$ lift to cycles of G_r. Walks at x_0 that stem from such cycles, those that generate $\pi_1^\varrho(G, x_0)$, therefore lie in the characteristic subgroup of p_r. By definition of p_r, this is the group $\pi_1^r(G, x_0)$.

(iii)→(iv). By definition, $\pi_1^r(G, x_0) \subseteq \pi_1^\varrho(G, x_0)$ for $r \leqslant \varrho$, so (iii) implies that these groups coincide. This implies (iv) by [H02, 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 coverings, the same is true for lifts of the same vertex of G to G_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 the implication (iii) $\to$(i) in Lemma 2.5.2, it suffices to show that every closed walk W in G based at x_0 that stems from a cycle O of length at most $\varrho := \ell - 1$ 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 $r \geqslant 2$. Given a connected graph G, 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.*

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 $r \geqslant 2$). By possible choosing $W \in \mathcal{W}_r$ such that it connects the same endvertices but n is minimum, we may assume that W repeats no vertices in X by definition of $\mathcal{W}_r(X)$, so 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. We also keep * at each reference to a statement assuming it.

Lemma 2.5.5. *Let G be connected and $r \leqslant 2$. Then every tight r-local separator of G has size 1.*

Proof. Since G is simple, G contains no cycles of length $\leqslant 2$. 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. □

An r-tomic set X is *r-locally 1-covered* if 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. Recall from Section 2.4.2 that the deck transformations γ of p_r act on the set of r-toms of $p_r^{-1}(X)$ by $Y \mapsto \gamma(Y)$. Thus, an r-tomic set is r-locally 1-covered if and only if we replace 'some' by 'every' in the definition of r-locally 1-covered.

Observation 2.5.6. *If $X \subseteq V(G)$ is r-locally 1-covered, then the action of the group $\Gamma_r(G)$ of deck transformations of p_r on the r-toms of $p_r^{-1}(X)$ is transitive and fixed-point free, i.e. the identity id_{G_r} is the only deck transformation that maps any r-tom on itself.*

Lemma 2.5.7 *. *Assume that $kr/2 < \Delta_r(G)$ for a connected graph G and $k \in \mathbb{N}$. Then every r-tomic vertex set X of size at most k is r-locally 1-covered. Moreover, the r-toms of $p_r^{-1}(X)$ are precisely the r-locally closed lifts of X.*

Proof. Let $\hat{X}$ be an arbitrary r-tom of $p_r^{-1}(X)$. As X is r-tomic, p_r maps $\hat{X}$ onto X by Projection 2.4.10. 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 $\hat{W}$ so 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 our assumption of $(*)$ implies $k = 1$ and $r \leqslant 1$. Then the $\hat{x}_i$ have distance 1 but are contained in the same fibre of p_r. This contradicts that G has no loops.

Moreover, since X is r-locally 1-covered, its r-toms are r-locally closed lifts of X. Conversely, every r-locally closed lift $\hat{X}$ of X contains an r-tom $\hat{Y}$ of $p_r^{-1}(X)$. As X is r-tomic, $\hat{Y}$ meets $p_r^{-1}(x)$ for every $x \in X$ by Projection 2.4.10, and huts $\hat{X} = \hat{Y}$. □

Example 2.5.8. Lemma 2.5.7 * fails if we do not assume $(*)$: Consider as G a cycle of length k, and the 1-tomic set $X = V(G)$, then $p_1^{-1}(X) = V(G_1)$ is 1-tomic but has infinite size, as G_1 is a double ray.

Projection 2.5.9 *. *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-tomic and r-locally 1-covered, and $\hat{X}$ is an r-tom of $p_r^{-1}(X)$. In particular, $\hat{X}$ is an r-locally closed lift of X, and so the projection $s := p_r(\hat{s})$ is 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-covered by Lemma 2.5.7 *. Since $\hat{X}$ is r-tomic, it is contained in some r-tom $\hat{Y}$ of $p_r^{-1}(X)$. Since X is r-locally 1-covered, p_r restricts to a bijection $\hat{Y} \to X$. As p_r restricted to $\hat{X} \subseteq \hat{Y}$ is already onto X, we have $\hat{X} = \hat{Y}$, so $\hat{X}$ is an r-tom of $p_r^{-1}(X)$. □

Lift 2.5.10 *. *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-covered and s can be lifted. In particular, the lifts of s to G_r form a Γ_r-orbit of the set of r-local separations of G_r.*

Proof. The separator X is r-locally 1-covered by Lemma 2.5.7 *. Hence, X lifts to an r-tom $\hat{X}$ of $p_r^{-1}(X)$, which thus is r-locally closed. By Lift 2.4.19, s lifts to an r-local separation of G_r with separator $\hat{X}$. The 'in particular'-part follows with Observation 2.5.6. □

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 (with parts G_h) 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 for now that the adhesion graphs of $(H, \mathcal{G})$ are pairwise disjoint and non-empty. For every edge $e = hh'$ of H we define $s^e := \{E_h^e, X^e, E_{h'}^e\}$, where $X^e := V(G_e)$,

$E^e_h := \partial_{G_h} X^e$ and $E^e_{h'} := \partial_{G_{h'}} X^e$. The axioms for graph-decompositions are easily seen to imply that, since our adhesion graphs are disjoint, the edge sets E^e_h and $E^e_{h'}$ form a near-partition of $\partial_G X^e$.

In this context, a *traversal* of s^e is a walk in G whose first edge is in E^e_h, whose last edge is in $E^e_{h'}$, and all whose internal vertices are in X^e. A closed walk $W = v_1 e_1 v_2 \dots v_n e_n v_1$ 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 subgraph of H formed by all the edges $e \in E(H)$ (and their endvertices) for which W traverses s^e.

Example 2.5.11. Assume that G is a connected graph. Let $(H, \mathcal{G})$ be a graph-decomposition (with parts G_h) of G such that

(i) every part G_h is complete, and

(ii) the adhesion graphs G_e are disjoint.

Then all induced cycles of G of length > 3 have length $\geqslant g(H)$. By Lemma 2.5.3, therefore, $\Delta_r(G) \geqslant g(H)$ for every $r \geqslant 3$.

Proof. Let O be an induced cycle in G. In what follows, we interpret O also as a closed walk once around O. If H^O is empty then, by the definition of graph-decompositions, 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, the intersection $O \cap G_e$ thus 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 of O in G_{h_i} are distinct for different i, by (ii). Hence, O has at least $n \geqslant g(H)$ edges. □

Let $(H, \mathcal{G})$ be a graph-decomposition (with parts G_h) of G, and assume that its 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 an initial segment equals k.

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. Let $W = v_1 e_1 v_2 \dots v_n e_n v_1$ be a closed walk in G. If $v_i e_i \dots e_k v_{k+1}$ is a subwalk of (a cyclic shift of) W that is a traversal of s^f for some $f \in E(H)$, then we refer to the indices j with $i < j < k$ as *bad indices* of the walk W. All indices of W that are not bad are *good*. We denote by $f(W)$ the number of all good indices of W. We then let $\ell(H, \mathcal{G})$ be the minimum $f(W)$ among all closed walks W in G which traverse at least one s^e and which traverse s^e at most once for each edge $e \in H$. 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 distinct edges e, f of H at h. If no such distinct edges e, f exist, we let $d(H, \mathcal{G}) := \infty$.

Corollary 2.5.12. *Let $(H, \mathcal{G})$ be a graph-decomposition (with parts G_h) 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; therefore, we can assume that there exists a closed walk W in G that witnesses $\ell(H,\mathcal{G}) < \infty$. 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. As G is finite and the adhesion graphs are disjoint, H and H^W are also finite. Hence, H^W contains a cycle O.

Let h be a node of O. Let e, f be the two edges on O that are incident to h. The walk W contains a subwalk W_h in G_h between two distinct adhesion graphs G_e, G_f of $(H,\mathcal{G})$, as W traverses s^e, s^f at most once. We can and do choose W_h so that its internal vertices are not contained in the adhesion graphs G_e, G_f. Then W_h has length at least $d(H,\mathcal{G})$. Since h has degree 2 in O, all indices of W that appear at the edges of W_h are good indices. As we find W_h with these properties for each node h of O, we find that $f(W) \geqslant g(H)\cdot d(H,\mathcal{G})$. □

Example 2.5.13. Let $(H,\mathcal{G})$ be a graph-decomposition (with parts G_h) of a finite connected graph G and an integer $r \geqslant 0$ such that

(i) every part G_h is connected and $\pi_1(G_h) = \pi_1^r(G_h)$,

(ii) its adhesion graphs G_e are 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 Corollary 2.5.12. The disjointness of the adhesion graphs given by (ii) ensures 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 say that a closed walk W in G 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 $\Delta_r(G) \geqslant \ell(H,\mathcal{G})$, it suffices by Lemma 2.5.2 (iii)→(i) 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

$$\textit{every closed walk } W \textit{ in } G \textit{ with } f(W) < \ell(H,\mathcal{G}) \textit{ is generated by short cycles.} \tag{2.5}$$

We prove (2.5) by induction on the number of subwalks of (cyclic shifts of) W that are traversals of some s^e (for $e \in E(H)$).

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). If W traverses at least one s^e and W traverses s^e at most once for every edge e of H, then W is itself a candidate in the definition of $\ell(H,\mathcal{G})$ and hence $f(W) \geqslant \ell(H,\mathcal{G})$, which contradicts that W has length $< \ell(H,\mathcal{G})$ by assumption.

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 $W_2 := v_j P^- v_i W v_j$. As W is homotopic to $W_1(v_1 W^- v_j)W_2(v_j W v_1)$, it suffices to show that W_1 and W_2 are generated by short cycles. Now both W_1 and W_2 have less traversals than W: By construction, each traversal of some s^f of $(H, \mathcal{G})$ in W_1 or W_2 is also a traversal in W, except possibly for the new ones using P or P^-, of which we may have created at most one in W_1 and in W_2 (but these need not be traversals at all). While W has the two traversals P_i and P_j of s^e, W_1 and W_2 contain neither of these two and each possibly a single new one. Thus, W_1 and W_2 each have at least one traversal less than W, and we can hence apply the induction hypothesis to both, completing the proof. □

2.5.3 Local separations whose separator is a clique

Recall from the example of Figure 2.17 that an r-tomic separator $\hat{X}$ of G_r can be much larger than the r-local separator X of G to which it projects. This cannot happen, however, if $\hat{X}$ induces a clique – an important case applied throughout [AKK25]:

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.*

For the proof of Lift 2.5.14, we need a lemma:

Lemma 2.5.15. *Let G be a connected graph and $r \geqslant 3$ an integer. Then:*

(i) *For every clique $\hat{X} \subseteq V(G_r)$, p_r is injective on $\hat{X}$, and its projection $p_r(\hat{X})$ is a clique of G.*

(ii) *For every clique $X \subseteq V(G)$ and every lift $\hat{x}$ of some vertex $x \in X$ to G_r, there exists a unique clique $\hat{X}$ of G_r containing $\hat{x}$ such that p_r restricts to a bijection from $\hat{X}$ to X.*

(iii) *$(\hat{X} \cup N(\hat{X})) \cap (\gamma(\hat{X}) \cup N(\gamma(\hat{X}))) = \emptyset$ for every clique $\hat{X} \subseteq V(G_r)$ and every $\gamma \in \Gamma(p_r) \setminus \{\mathrm{id}_{G_r}\}$.*

Moreover, for every clique $X \subseteq V(G)$, the r-toms of $p_r^{-1}(X)$ in G_r are precisely the cliques $\hat{X} \subseteq V(G_r)$ for which p_r restricts to a bijection from $\hat{X}$ to X.

Proof of Lemma 2.5.15. (i) holds by the definition of covering. (ii) follows from $r \geqslant 3$ and the uniqueness of path-lifting in coverings.

For (iii), suppose for a contradiction that there are $\gamma \in \Gamma(p_r) \setminus \{\mathrm{id}_{G_r}\}$ and a clique $\hat{X} \subseteq V(G_r)$ such that $\hat{X} \cup N(\hat{X})$ and $\gamma(\hat{X}) \cup N(\gamma(\hat{X}))$ share a vertex $\hat{y}$. Then there are a vertex $\hat{x}_0 \in \hat{X}$ and a vertex $\hat{x}_1 \in \gamma(\hat{X})$ both with distance $\leqslant 1$ to $\hat{y}$ in G_r. As $\hat{x}_0' := \gamma(\hat{x}_0)$ has distance $\leqslant 1$ from $\hat{x}_1 \in \gamma(\hat{X})$, the two vertices $\hat{x}_0$ and $\hat{x}_0'$ in the same fibre of p_r thus have distance $\leqslant 3 \leqslant r$. This contradicts Lemma 2.5.1.

For the 'moreover'-part, note first that any lift $\hat{X}$ of X to G_r that is a clique must be contained in some r-tom of $p_r^{-1}(X)$. For the converse inclusion, consider any r-tom $\hat{Y}$ of $p_r^{-1}(X)$. By (ii), $\hat{Y}$ contains a lift $\hat{X}$ of X that is a clique. Since $\Gamma(p_r)$ acts transitively on the fibres of p_r^{-1}, (iii) implies that $\hat{X}$ cannot be joined by an edge to any vertex in $p_r^{-1}(X) \setminus \hat{X}$. Now if $\hat{Y} \setminus \hat{X}$ is non-empty, then there is a short cycle in G_r meeting both $\hat{X}$ and a vertex $\hat{y} \in \hat{Y} \setminus \hat{X}$. Thus, $\hat{y}$ has distance at most $r-1$ to some vertex in $\hat{X}$, and

since $\hat{X}$ is a clique, $\hat{y}$ has distance at most r to every vertex in $\hat{X}$. But $\hat{X}$ is a lift of X, and so $\hat{X}$ contains a vertex $\hat{y}'$ that lies in the same fibre as $\hat{y}$. 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. By the 'moreover'-part of Lemma 2.5.15, the lifts $\hat{X}$ of X to G_r that are cliques are precisely the r-toms of $p_r^{-1}(X)$ in G_r; in particular, these lifts $\hat{X}$ are precisely the r-locally closed lifts of X. Thus, the statement follows immediately from Projection 2.4.18 and Lift 2.4.19. □

2.6 Links and crossing

As the next step regarding r-local separations, we develop notions of *crossing* and *nestedness* 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 (Lift 2.6.6 * and Lift 2.6.7 *). Moreover, if s and t cross, then every lift of s crosses exactly one lift of t (Lift 2.6.8 *).

2.6.1 From global to local, and correspondence in the local cover

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 of x, y, x', y'. We emphasise that O may intersect $X \cup Y$ in more than four vertices. This definition is inspired by [C22].

Two vertex sets $X, Y \subseteq V(G)$ are *r-coupled* in G if either X and Y share a vertex or there is a short cycle in G that alternates between X and Y. Further, let $X \subseteq V(G)$ and let $\{F_1, Y, F_2\}$ be any r-local separation of G. The *X-link for F_i in $\{F_1, Y, F_2\}$* 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.

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 would be empty, and this would contradict 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 an 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)**.** *Let $\{A_1, A_2\}$ and $\{C_1, C_2\}$ be separations of a connected graph 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 W in Y. Since x lies in the X-link for C_1, all vertices of W that precede y are contained in $C_1 \smallsetminus C_2$. Hence, y is preceded on W by an edge in F_1. 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)**.** *Let s and t be separations of a connected graph 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. *Whenever G is a connected graph, $(k-1)r < \Delta_r(G)$ implies $kr/2 < \Delta_r(G)$, for all integers $k, r \geqslant 0$.*

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 $(k-1)r - kr/2 = (k/2 - 1)r > 0$, so $kr/2 < (k-1)r$ and the desired implication follows. □

Recall that we assume that $r \geqslant 2$ or $k = 1$, if we add * to a statement. We assume * for the rest of Section 2.6.

The next lemma addresses lifting the notion of 'r-coupled'. For this, recall that for an r-tomic vertex set X in a connected graph G, the r-toms $\hat{X}$ of $p_r^{-1}(X)$ are the lifts of X by Lemma 2.6.5 and Lemma 2.5.7*.

Lift 2.6.6*. *Let X and Y be r-coupled, r-tomic vertex sets of size $\leqslant k$ in a connected graph G such that $(k-1)r < \Delta_r(G)$. 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 $\hat{X} \cup V(\hat{H})$. Any $v \in X \cap Y$ then has unique lift $\hat{v}$ in $\hat{X} \cap \hat{Y}$, and $\hat{Y}$ is the unique r-tom of $p_r^{-1}(Y)$ that is r-coupled with $\hat{X}$. Every short cycle O in G that alternates between X and Y then has a unique lift $\hat{O}$ that alternates between $\hat{X}$ and $\hat{Y}$.*

Proof. First, assume that $r \leqslant 1$. Then $k = 1$ by *, so $|X| = |Y| = 1$. Since there are no cycles of length at most 1 and as X and Y are r-coupled, we thus have $H = \emptyset$ and $X = Y$, say $X = Y = \{z\}$. As G has no loops (at z), $p_r^{-1}(z)$ is an independent set vertex set of G_r, whose r-toms are formed by its single vertices. Every r-tom $\hat{X} = \{\hat{z}\}$ is the unique r-tom $\hat{Y}$ of $p_r^{-1}(Y)$ that meets $\hat{X}$. Hence, we may assume that $r \geqslant 2$.

Existence of $\hat{Y}$. Let an r-tom $\hat{X}$ of $p_r^{-1}(X)$ be given. Since X is r-locally 1-covered, by Lemma 2.6.5 and Lemma 2.5.7*, $\hat{X}$ is a lift of X. As 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 $v \in X \cap Y$ exists, let $\hat{v}$ be its unique lift in $\hat{X}$. If O exists, let $\hat{O}$ denote any lift of O that meets $\hat{X}$. As $\hat{O}$ is short, it meets no other r-tom of $p_r^{-1}(X)$. Similarly, let $\hat{Y}$ denote the unique r-tom of $p_r^{-1}(Y)$ that contains $\hat{v}$ if v exists. if not, let $\hat{Y}$ be the unique r-tom of $p_r^{-1}(Y)$ that meets $\hat{O}$. As earlier, $\hat{Y}$ is a lift of Y. 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 $\hat{X} \cup V(\hat{H})$, as desired. Finally, since $\hat{X}$ is a lift of X and also r-locally closed as an r-tom of $p_r^{-1}(X)$, the short cycle $\hat{O}$ is the unique lift of O that meets $\hat{X}$; in parituclar, it is the unique lift of O alternating between $\hat{X}$ and $\hat{Y}$.

Uniqueness of $\hat{Y}$. 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 $\hat{y}_1, \hat{y}_2$, respectively. If $\hat{y}_i \in V(\hat{H}) \setminus \hat{X}$, then let $\hat{O}_i$ be a short cycle in G_r that contains $\hat{y}_i$ and intersects $\hat{X}$ in some two vertices $\hat{x}_i, \hat{x}_i'$. Further, let $\hat{y}_i'$ denote the unique vertex in $\hat{Y}_{3-i}$ with $p_r(\hat{y}_i') = p_r(\hat{y}_i)$.

Suppose first that both $\hat{y}_i$ lie in $\hat{X}$. Since $\hat{X}$ and $\hat{Y}_2$ are r-tomic, there are an $\hat{y}_1$–$\hat{y}_2$ path and an $\hat{y}_2$–$\hat{y}_1'$ path in G_r, both of length $\leqslant (k-1)r/2$, by Lemma 2.5.4. So $\hat{y}_1$ and $\hat{y}_1'$ have distance $\leqslant (k-1)r$, contradicting our assumption that $(k-1)r < \Delta_r(G)$.

Suppose second that $\hat{y}_1$ lies in $\hat{X}$ and that $\hat{y}_2$ lies on $\hat{O}_2$. Since $\hat{X}$ is r-tomic, there is a $\hat{y}_1$–$\{\hat{x}_2, \hat{x}_2'\}$ walk $Q(\hat{X}) \in \mathcal{W}_r(\hat{X})$. By choosing $Q(\hat{X})$ of minimum length, we may assume that $Q(\hat{X})$ ends in $\hat{x}_2$ and avoids $\hat{x}_2'$, say. Since $Q(\hat{X})$ avoids $\hat{x}_2'$, there is a $\hat{y}_1$–$\hat{x}_2$ 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 $\leqslant r/2$ in $\hat{O}_2$ from $\hat{x}_2$ to $\hat{y}_2$. Since $\hat{Y}_2$ is r-tomic, there is an $\hat{y}_2$–$\hat{y}_1'$ 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 $\hat{y}_1$ to $\hat{y}_1'$ and has length $\leqslant (k-1)r$, contradicting our assumption that $(k-1)r < \Delta_r(G)$.

The case that $\hat{y}_1$ lies on $\hat{O}_1$ while $\hat{y}_2$ lies in $\hat{X}$ follows from the previous case by symmetry.

Suppose finally that both $\hat{y}_i$ lie on $\hat{O}_i$. Since $\hat{X}$ is r-tomic, there is a $\{\hat{x}_1, \hat{x}_1'\}$–$\{\hat{x}_2, \hat{x}_2'\}$ walk $Q(\hat{X}) \in \mathcal{W}_r(\hat{X})$; choose this of minimum length. Possibly by renaming some vertices, we may assume that $Q(\hat{X})$ links $\hat{x}_1$ to $\hat{x}_2$. If $\{\hat{x}_1, \hat{x}_1'\}$ and $\{\hat{x}_2, \hat{x}_2'\}$ 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 $\hat{x}_1$–$\hat{x}_2$ 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 $\hat{y}_i$–$\hat{x}_i$ 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 $\hat{y}_2$–$\hat{y}_1'$ 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 $\hat{y}_1$ to $\hat{y}_1'$ and has length $\leqslant (k-1)r$, contradicting our assumption 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 lift of t such that $\hat{X}$ and $\hat{Y}$ are r-coupled, which exists and is unique 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 2.5.10* 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, it has a unique lift $\hat{v}$ in $\hat{X} \cap \hat{Y}$ by Lift 2.6.6*. If O exists, it has a unique lift $\hat{O}$ that alternates between $\hat{X}$ and $\hat{Y}$, again by Lift 2.6.6*.

Since $\hat{s}$ is a lift of s, its separator $\hat{X}$ is a lift of X and r-locally closed. Further, by Lemma 2.5.7*, $\hat{X}$ is an r-tom of $p_r^{-1}(X)$. Analogously, $\hat{Y}$ is an r-locally closed lift of Y and thus an r-tom of $p_r^{-1}(Y)$. By Lift 2.6.6*, $\hat{Y}$ is the unique r-tom of $p_r^{-1}(Y)$ which is r-coupled with $\hat{X}$. Thus, (iii) holds, and this immediately yields (iv).

We now claim that

$$\text{the lift of the } X\text{-link for } F_1 \text{ that lies in } \hat{X} \text{ is included in the } \hat{X}\text{-link for } \hat{F}_1. \tag{2.6}$$

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 is a concatenation of r-local X-paths $P_1, \dots, P_n$ starting at x, 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 $\hat{x}$. By the path-lifting property of coverings, this defines lifts $\hat{P}_i$ of all the P_i and a lift $\hat{e}$ of e. As $\hat{X}$ is r-locally closed, the $\hat{P}_i$

are r-local $\hat{X}$-paths. Hence, $\hat{W} \in \mathcal{W}_r(\hat{X})$. It is immediate that $\hat{e}$ is the first edge of $\hat{W}$ in $\partial p_r^{-1}(Y)$. Thus, 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 $\hat{P}_i$, it either lies on a short cycle in G_r that has two vertices in $\hat{X}$ or is an $\hat{X}$-edge. Thus, $\hat{X}$ and $\hat{Z}$ are r-coupled. Since $\hat{Y}$ is the unique r-tom of $p_r^{-1}(Y)$ that is r-coupled with $\hat{X}$, Lift 2.6.6* yields $\hat{Z} = \hat{Y}$. Therefore, $\hat{e} \in \partial Y$. Then $\hat{e} \in \hat{F}_1$ as well. This completes the proof of (2.6).

Next, we claim that

$$\text{the projection of the } \hat{X}\text{-link for } \hat{F}_1 \text{ is included in the } X\text{-link for } F_1. \tag{2.7}$$

So let $\hat{x}$ be a vertex in the $\hat{X}$-link for $\hat{F}_1$, and let $\hat{W} \in \mathcal{W}_r(\hat{X})$ starting at $\hat{x}$ whose first edge $\hat{e}$ in $\partial\hat{Y}$ lies in $\hat{F}_1$. If $e := p_r(\hat{e})$ is also first edge f of $W := p_r(\hat{W})$ in ∂Y, then the projection $x := p_r(\hat{x})$ is in the X-link for F_1 and we are done.

So let f be the first edge of W which is in ∂Y. Let $\hat{f}$ be its lift in $\hat{W}$, and let $\hat{Z}$ be the r-tom of $p_r^{-1}(Y)$ whose $\partial\hat{Z}$ contains $\hat{f}$. Since $\hat{W} \in \mathcal{W}_r(\hat{X})$, $\hat{f}$ either lies on a short cycle that meets $\hat{X}$ in at least two vertices or is an $\hat{X}$-edge. Since $\hat{Y}$ is the unique r-tom of $p_r^{-1}(Y)$ that is r-coupled with $\hat{X}$, Lift 2.6.6* yields $\hat{Z} = \hat{Y}$. In particular, $\hat{f} = \hat{e}$, as desired.

Both (2.6) and (2.7) extend to all four links by analogue arguments, so (i) and (ii) hold. The 'moreover'-part is immediate from (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) *every lift of s crosses exactly one lift of t.*

In the context of Lift 2.6.8*, we recall that s and t can be lifted by Lift 2.5.10* and Lemma 2.6.5.

Proof. (i)→(ii). This follows directly from the 'moreover'-part of Lift 2.6.7*. (Note that (i) implies, by definition of 'cross', that s and t are r-coupled, which Lift 2.6.7* requires.)

(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 of p_r on the r-locally closed lifts of X (equivalently by Lemma 2.5.7*: r-toms of $p_r^{-1}(X)$), 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. Since p_r isomorphically maps short cycles of G_r to short cycles in G, X and Y are r-coupled too. Thus, Lift 2.6.6* yields that $\hat{t}$ is the only lift of t that crosses $\hat{s}$, as desired.

(iii)→(i). By Lift 2.5.10* and Lemma 2.6.5, s and t have lifts. By (iii), some of these cross, say $\hat{s}$ and $\hat{t}$ with separators $\hat{X}$ and $\hat{Y}$. As in the proof of (ii)→(iii), X and Y are r-coupled. By the 'moreover'-part of Lift 2.6.7*, s and t cross, as desired. □

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-coupled, r-tomic separators X and Y of size $\leqslant k$, respectively, and assume that $(k-1)r < \Delta_r(G)$. 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 2.5.10 * 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 the fact 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 shall follow a recurring scheme: we define the local notions, show that they agree with the respective global counterpart on r-local covers, and prove lifting/projection lemmas under a suitable assumption on the displacement of p_r. We work along this three-step process first for $\top$-stars (Correspondence 2.7.3, Projection 2.7.11 * and Lift 2.7.12 *), then for 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 that represent the r-local bottlenecks (Construction 2.7.16, Correspondence 2.7.17, Theorem 2.7.18 * and Theorem 2.7.19 *).

2.7.1 Local and global $\top$-stars in the local cover

We need a way to localise $\top$-stars. To do this, let us first characterise $\top$-stars in more local terms than by their definition:

Lemma 2.7.1. *A set $\sigma = \{ (A_i, B_i) : i \in [3] \}$ of three distinct oriented separations (A_i, B_i) of a graph G with separators X_i is a $\top$-star if and only if*

(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.*

Proof. This is immediate from the definition of $\top$-star. □

An *r-local* $\top$*-star* in G is a set $\sigma = \{ (E_i, X_i, F_i) : i \in [3] \}$ of three oriented r-local separations such that

(LY1) the E_i 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 r-local 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 $X_{ij} = X \setminus X_k$ whenever $\{i, j, k\} = [3]$. 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 2.7.2. *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 definition (Lemma 2.4.2). 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. □

Given a set σ of (possibly oriented) separations of G, we write $\sigma_r := \{s_r : s \in \sigma\}$ for the 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. Then $(\sigma, s) \mapsto (\sigma_r, s_r)$ 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. Recall that, by Lemma 2.7.1, $\top$-stars are characterised by (Y1) and (Y2).

Well-defined. Let $\sigma = \{ (A_i, B_i) : i \in [3] \}$ be a relevant $\top$-star in G with base $\{A_1, B_1\}$. For every i, let (E_i, X_i, F_i) denote the r-local separation induced by (A_i, B_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 2.4.5.

Let us now show that $\sigma_r = \{ (E_i, X_i, F_i) : i \in [3] \}$ is an r-local $\top$-star. (LY1) and (LY2) follows immediately from (Y1) and (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 2.4.4. As the (A_i, B_i) are separations of G which satisfy (Y1), every component of $G - X$, in particular K, has its vertices in some $A_i \setminus B_i$. Thus, $F \subseteq E(X, A_i \setminus B_i) \subseteq E_i$ 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 from Correspondence 2.4.8 and Correspondence 2.4.4 (applied to the r-local separations $\{E_i, X, \partial X \smallsetminus E_i\}$ for $i = 2, 3$) that σ_r is relevant with base $\{E_1, X_1, F_1\}$.

Injective. Since G is connected, this follows from the definition of r-local separations induced by separation, i.e. the map $s \mapsto s_r$.

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, let us find separations (A_i, B_i) of G that induce the (E_i, X_i, F_i). Since $X := \bigcup_{i\in[3]} X_i$ is r-tomic, by Lemma 2.7.2, the r-local separations $(E_i, X, \partial X \smallsetminus E_i)$ are induced by separations (A'_i, B'_i) of G, by Correspondence 2.4.6, so, $E(X, A'_i \smallsetminus B'_i) = E_i$. Now let $A_i := (A'_i \smallsetminus B'_i) \cup X_i$ and $B_i := B'_i$. Then $A_i \smallsetminus B_i = A'_i \smallsetminus B'_i$ and $E_G(X_i, A_i \smallsetminus B_i) = E_i$. Thus, $E(X_i, B_i \smallsetminus A_i) = \partial X_i \smallsetminus E_i = F_i$. 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 $(E_i, X, \partial X \smallsetminus E_i)$ are r-local separations whose sets E_i near-partition ∂X, by (LY1). Hence, for the corresponding separations (A'_i, B'_i) the sets $A'_i \smallsetminus B'_i$ near-partition the components of $G - X$, by (LY3) and Correspondence 2.4.4. Since $A'_i \smallsetminus B'_i = A_i \smallsetminus B_i$, the sets $A_i \smallsetminus B_i$ near-partition $V(G) \smallsetminus X$, giving (Y1). As (LY2) is identical to (Y2), Lemma 2.7.1 implies that σ is a $\top$-star.

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 2.4.8 and Correspondence 2.4.4 (applied to the r-local seps $\{E_i, X, \partial X \smallsetminus E_i\}$ for $i = 2, 3$). 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-covered, but which also satisfies $f(k) \leqslant 2(k-1)$ for every positive $k \in \mathbb{N}$:

$$f(k) := \lfloor 3k/2 \rfloor \text{ for } k \geqslant 3 \text{ and } f(k) := 0 \text{ for } k \leqslant 2. \tag{2.8}$$

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 $kr/2 \leqslant r < \Delta_r(G)$, by Lemma 2.5.1, while $k \geqslant 3$ gives $\Delta_r(G) > f(k) \cdot r/2 = \lfloor 3k/2 \rfloor \cdot r/2 \geqslant kr/2$. □

The following Lemma 2.7.6 is of the same flavour as [C22, 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 an X-edge.*

Proof. Let $\{x, y\} := X$. Let F_i be a tight component in E_i for $i = 1, 2$. For each i, pick a short cycle $O_i \subseteq G$ that includes an X-path starting and ending with edges in F_i. 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 the union of the desired paths, and similarly for O_2. Hence, we may assume that both X-paths in O_i start and end with 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 at most 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 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 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 our assumption that $|X| \geqslant 3$. Thus, all the X_i are distinct, and by (LY2) every two X_i share an element. Hence, the centre of the $\top$-star is empty and every ij-link consists of a unique vertex, which we denote by x_{ij}. Thus, $X = \{x_{12}, x_{23}, x_{13}\}$.

Since $\{E_1, X_1, F_1\}$ is tight, we may use Lemma 2.7.6 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 single edge $x_{12}x_{13}$, 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 the other endvertex of f is x_{23}, then O contains X as desired. If not, then $f \in \partial X \setminus \partial X_3$. Since $f \in F_1$, we have $f \in \partial X \setminus (E_1 \cup E_3)$ and hence $f \in E_2$ by (LY1). Let us show that $e \in F_2$. If $e = x_{12}x_{13}$, then $e \in \partial X_2 \setminus \partial X$, so $e \in F_2$ by (LY1). Otherwise, e lies in $E_1 \subseteq \partial X$ as well as 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 $\{E_2, X_2, F_2\}$. As O is a short cycle containing 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 $|X| \leqslant 2$ or X is contained in a short cycle, then X is r-locally 1-covered.*

Proof. If $|X| \leqslant 2$, then $\Delta_r(G) > r \geqslant k \cdot r/2$ for $k \leqslant 2$, Lemma 2.5.1, so X is r-locally 1-covered by Lemma 2.5.7*.

Now assume that X is contained in some short cycle $O \subseteq G$. Let $\hat{O}$ be a lift of O. Then $\hat{O}$ is a short cycle too, so $\hat{X} := V(\hat{O}) \cap p_r^{-1}(X)$ is contained in an r-tom of $p_r^{-1}(X)$. Let us show that this consists of $\hat{X}$ only.

If not, there exists a short cycle which joins some $\hat{x} \in \hat{X}$ to some $\hat{y} \in p_r^{-1}(X) \setminus \hat{X}$. 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

$$d_{G_r}(\hat{y}, \hat{y}') \leqslant d_{G_r}(\hat{y}, \hat{x}) + d_{G_r}(\hat{x}, \hat{y}') \leqslant r/2 + r/2 = r,$$

which 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-covered.*

Proof. The set X is r-tomic by Lemma 2.7.2. If $k \leqslant 2$, then X is r-locally 1-covered by Lemma 2.7.7 and Lemma 2.7.8. So we may assume that $k \geqslant 3$. By (LY2), X has size $\leqslant \lfloor 3k/2 \rfloor$. Then X is r-locally 1-covered by Lemma 2.5.7 *, as $\lfloor 3k/2 \rfloor \cdot r/2 = f(k) \cdot r/2 < \Delta_r(G)$. □

Lemma 2.7.10 *. *Let G be a connected graph such that $f(k) \cdot r/2 < \Delta_r(G)$. Let $\hat{\sigma}$ be a relevant r-local $\top$-star of order $\leqslant k$ in G_r. 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-covered. Moreover, $\hat{X}$ is an r-tom of $p_r^{-1}(X)$, and thus a lift of X.*

Proof. The set $\hat{X}$ is r-tomic in G_r by Lemma 2.7.2. 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-covered by Lemma 2.7.8.

So we may assume $k \geqslant 3$. By (LY2), $\hat{X}$ has size $\leqslant \lfloor 3k/2 \rfloor$, hence so does its projection X. Thus, X is r-locally 1-covered 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 an r-tom of $p_r^{-1}(X)$, since the r-tomic set $\hat{X}$ is contained in an r-tom of $p_r^{-1}(X)$ r-tomic and its projection $X = p_r(\hat{X})$ is r-locally 1-covered. By the definition of r-locally 1-covered, $\hat{X}$ is a lift of X. □

Projection 2.7.11 * (Local $\top$-star). *Let G be a connected graph such that $f(k) \cdot r/2 < \Delta_r(G)$. Let $\hat{\sigma}$ be a relevant r-local $\top$-star in G_r of order $\leqslant k$ with base $\hat{s}$. 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-covered and $\hat{X}$ is an r-tom of $p_r^{-1}(X)$, and thus a lift of X. Hence, the subsets $\hat{X}_i$ of $\hat{X}$ are r-locally closed lifts of their projections X_i. Thus, by Projection 2.4.18, each $\hat{s}_i$ projects to an r-local separation $s_i = (E_i, X_i, F_i)$ of G, which is tight for $i = 1$.

Since $\hat{X}$ is a lift of X, (LY1) and (LY2) for $p_r(\hat{\sigma}) = \{ s_i : i \in [3] \}$ follow immediately from (LY1) and (LY2) of $\hat{\sigma}$. (LY3) and the relevance of $\hat{\sigma}$ with base $\hat{s}_1$ imply the same for σ and s_1 by Lift 2.4.16. □

Lift 2.7.12 * (Local $\top$-star). *Let G be a connected graph such that $f(k) \cdot r/2 < \Delta_r(G)$. 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 $\{E_1, X_1, F_1\}$. Let $\hat{X}$ be any 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 2.4.19, 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-covered by Lemma 2.7.9 *, so $\hat{X}$ is a lift of X. Thus, each X_i has a unique lift $\hat{X}_i \subseteq \hat{X}$, which is r-locally closed since $\hat{X}$ is r-locally closed as an r-tom of $p_r^{-1}(X)$. Hence the lifts $\hat{s}_i = (\hat{E}_i, \hat{X}_i, \hat{F}_i)$ defined as in Lift 2.4.19 exist and are unique, and $\{\hat{E}_1, \hat{X}_1, \hat{F}_1\}$ is tight. As $\hat{X}$ is a lift of X, (LY1) and (LY2) transfer

from σ to $\hat{\sigma} := \{ \hat{s}_i : i \in [3] \}$. Lift 2.4.16 ensures that (LY3) and relevance transfer 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 satisfying:

(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.

We note that since the r-local separations in an r-local bottleneck are tight, their separators are r-tomic by Lemma 2.4.2.

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 $\beta \mapsto \beta_r := \{s_r : s \in \beta\}$ given by Correspondence 2.4.8 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_r satisfies (B_r) if and only if β satisfies (B). □

Projection 2.7.14 * (Local bottlenecks). *Let G be a connected graph such that $f(k) \cdot r/2 < \Delta_r(G)$. Let $\hat{\beta}$ be an r-local k-bottleneck in G_r. Then $p_r(\hat{\beta})$ is an r-local k-bottleneck in G.*

Proof. Hence, by Projection 2.5.9 * and Lemma 2.7.5, we may use Projection 2.4.18 to find that $\beta := p_r(\hat{\beta})$ is 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 $\hat{s}$. By Lift 2.7.12 *, the elements of σ have unique lifts to G_r whose separators are included in $\hat{X}$, and these form a relevant r-local $\top$-star $\hat{\sigma}$ in G_r of order $\leqslant k$ with base $\hat{s}$. 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 the elements of β can be lifted to G_r and the set $\hat{\beta}$ of all such lifts is an r-local k-bottleneck in G_r.*

Proof. By Lift 2.5.10 * and Lemma 2.7.5, the elements of β have r-locally 1-covered separators and lift to G_r. All these lifts, the elements of $\hat{\beta}$, are tight r-local k-separations of G_r by Lift 2.4.19; 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 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, and let $\mathcal{B}_r$ be any set of r-local bottlenecks. For every $k \in \mathbb{N}$, let $\mathcal{B}_r^k(G)$ be the set of the r-local k-bottlenecks in $\mathcal{B}$ and $S_r^k := \bigcup \mathcal{B}_r^k(G)$ for every $k \in \mathbb{N}$. Recursively for all $k \in \mathbb{N}$, define sets $N_r^k(G) \subseteq S_r^k$, as follows. Let $x_r^k \colon S_r^k \to \mathbb{N}$ assign to each r-local separation $s \in X_r^k$ the number of r-local separations in X_r^k that s crosses assuming for now that this number is finite. 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 $N_r^k(G) := \bigcup_{\beta \in \mathcal{B}_r^k} N_r^k(G, \beta)$.

Whenever we use the notation of $N_r^k(G)$ without reference to any set $\mathcal{B}_r$ of r-local bottlenecks, we set $\mathcal{B}_r$ to be the set of all r-local bottlenecks in G.

In full analogy to separations, we say that a set S of r-local separations of G is Γ-*canonical* for some subgroup Γ of $\mathrm{Aut}(G)$ if S is invariant under the action of Γ on the r-local separations of G, as given by $\{E_1, X, E_2\} \mapsto \{\varphi(E_1), \varphi(X), \varphi(E_2)\}$ for $\varphi \in \mathrm{Aut}(G)$. In the case of $\Gamma = \mathrm{Aut}(G)$, we just say that S is *canonical*. Further, a set $\mathcal{B}_r$ of r-local bottlenecks in G is Γ-*invariant* for a subgroup Γ of $\mathrm{Aut}(G)$ if the action of Γ on the r-local bottlenecks in G (which is inherited from the action of $\mathrm{Aut}(G)$ on the r-local separations of G) is invariant on $\mathcal{B}$.

Correspondence 2.7.17 (Nested set)**.** *Let G be a locally finite, connected graph, let $\mathcal{B}$ be a set of bottlenecks in G, and let $\mathcal{B}_r = \{\beta_r : \beta \in \mathcal{B}\}$ be the set of r-local bottlenecks given by Correspondence 2.7.13. 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 of the r-local separations s_r induced by the separations s in $N^k(G)$ given by Construction 2.3.6. In particular, $N_r^k(G)$ is a nested set of r-local separations. Moreover, if $\mathcal{B}_r$ is Γ-invariant for a subgroup Γ of* $\mathrm{Aut}(G)$*, then $N_r^k(G)$ is Γ-canonical.*

Proof. The map $\beta \mapsto \beta_r$ given by Correspondence 2.7.13 restricts for each $k \in \mathbb{N}$, to a bijection from the set $\mathcal{B}^k$ of all k-bottlenecks in $\mathcal{B}$ to the set $\mathcal{B}_r^k$ of the r-local k-bottlenecks in $\mathcal{B}_r$. 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 in a connected graph. Assume now that $k > 0$. Then Correspondence 2.4.8 together with the induction hypothesis yields 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 maps only to finite numbers. Moreover, $s \mapsto s_r$ restricts to a bijection $N^k(G, \beta) \to N_r^k(G, \beta_r)$ for every $\beta \in \bigcup \mathcal{B}^k$.

For the 'in particular'-part, recall that $N^k(G)$ is nested by Theorem 2.3.7, so Correspondence 2.6.4 implies that $N_r^k(G)$ is nested as well. The 'moreover'-part is immediate from the canonicity of Construction 2.7.16. □

Theorem 2.7.18* (Nested set). *Let G be a locally finite, connected graph such that $(k-1)r < \Delta_r(G)$. Let $\mathcal{B}_r$ be a set of r-local bottlenecks in G, and let $\hat{\mathcal{B}}_r$ be the set of r-local bottlenecks in G_r given by Lift 2.7.15* applied to $\mathcal{B}_r$. Then $N_r^k(G)$ 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 $N_r^k(G)$. In particular, $N_r^k(G)$ is nested. Moreover, if $\mathcal{B}_r$ is Γ-invariant for a subgroup Γ of $\mathrm{Aut}(G)$, then $N_r^k(G)$ is Γ-canonical.*

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 $\hat{\mathcal{B}}_r$ to the set $\mathcal{B}_r^j$ of all r-local j-bottlenecks in $\mathcal{B}_r$; the map is well-defined by Projection 2.7.14* and is surjective by Lift 2.7.15*. Moreover, Projection 2.7.14* and Lift 2.7.15* ensure that $\bigcup \hat{\mathcal{B}}_r^j$ is the set of all the lifts of 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 in a connected graph. 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*. Moreover, Lift 2.6.8* yields 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 a finite number, since the latter is, by Correspondence 2.7.17. Hence, $\hat{s} \mapsto p_r(\hat{s})$ 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 $N_r^k(G)$.

For the 'in particular'-part, recall that $N_r^k(G_r)$ is nested by Correspondence 2.7.17, so Lift 2.6.8* implies that $N_r^k(G)$ is nested as well. The 'moreover'-part is immediate from the canonicity of Construction 2.7.16. □

Theorem 2.7.19*. *Let G be a locally finite, connected graph such that $(k-1)r < \Delta_r(G)$. Let $\mathcal{B}_r$ be a set of r-local bottlenecks in G. Then $N_r^{\leqslant k}(G)$ meets every r-local bottleneck in $\mathcal{B}_r$ of order $\leqslant k$.*

Proof. Let $\hat{\mathcal{B}}_r = \{\hat{\beta}_r : \beta_r \in \mathcal{B}_r\}$ be the set of r-local bottlenecks in G_r given by Lift 2.7.15* applied to $\mathcal{B}_r$, and let $\hat{\mathcal{B}} = \{\hat{\beta} : \hat{\beta}_r \in \hat{\mathcal{B}}_r\}$ be the set of bottlenecks in G_r given by Correspondence 2.7.13. Let β_r be any r-local bottleneck in $\mathcal{B}_r$, and write $\hat{\beta}_r \in \hat{\mathcal{B}}_r$ and $\hat{\beta} \in \hat{\mathcal{B}}$ for the corresponding (r-local) bottlenecks in G_r. By Theorem 2.3.7 applied to $\hat{\mathcal{B}}$ in G_r, there is $\hat{s} \in \hat{\beta} \cap N^{\leqslant k}(G_r)$. So $\hat{s}_r \in \hat{\beta}_r \cap N_r^{\leqslant k}(G_r)$ by Correspondence 2.7.17 for $\hat{\mathcal{B}}_r$. Since $N_r^{\leqslant k}(G_r)$ for $\hat{\mathcal{B}}_r$ is equal to the set of all lifts of all elements of $N_r^{\leqslant k}(G)$ for $\mathcal{B}_r$ by Theorem 2.7.18*, we have $p_r(\hat{s}_r) \in \beta_r \cap N_r^{\leqslant k}(G)$. □

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 G at hand (see Section 2.7), it remains to reconstruct from just this

local information the graph-decompositions of G which previously, in [DJKK26], required full knowledge of the (usually infinite) cover G_r of G and its (also usually infinite) tree of tangles from Construction 2.3.6.

This begins with the challenge to construct a locally defined analogue of the tree-decomposition from Construction 2.2.3, which was based on a (global) tree of tangles of G_r. Building the tree-decomposition relied on the partial ordering of global separations, and it is not clear what a meaningful r-local version of this partial ordering could be.

To overcome this obstacle, we develop an alternative perspective on Construction 2.2.3 by introducing the notion 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); in particular, these cutouts give rise to the same tree-decomposition of G_r as the globally defined tree of tangles did. 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 any nested set N of r-local separations is isomorphic to that obtained by taking the set $\hat{N}$ of all the lifts to G_r of elements of N, applying our construction to $\hat{N}$ in G_r, 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 an oriented separation $s = (A_1, A_2)$ of G we denote its separator by $X = X(s)$ and define $L(s) := E_G(A_1 \setminus A_2, X) = E_G(A_1 \setminus A_2, A_2)$ and $R(s) := E_G(X, A_2 \setminus A_1) = E_G(A_1, A_2 \setminus A_1)$. We think of $L(s)$ and $R(s)$ as the *left* and *right* sides of the local separations induced by the separation s (whose names are purely based on the reading direction). Note that, for elements $s > s'$ of $\mathrm{ori}(N)$, any edge in $R(s) \cap R(s')$ has a vertex in $X(s) \cap X(s')$.

Let N be a nested set of separations of a graph G. We define a relation $\Sigma^{\vec{x}}_N$ on the set $\mathrm{ori}(N)$ of orientations of elements of N by letting

$$s_1 \Sigma^{\vec{x}}_N s_2 :\Leftrightarrow X(s_1) \text{ and } X(s_2) \text{ share a vertex } v \text{ and } s_1 > s_2^*, \\ \text{and there is no } s \in \mathrm{ori}(N) \text{ with } v \in X(s) \text{ and } s_1 > s > s_2^*.$$

Recall that s^* denotes the inverse (B, A) of a separation $s = (A, B)$. The relation $\Sigma^{\vec{x}}_N$ is symmetric. Recall from Construction 2.2.3 the definition of $\mathcal{T}(N)$. If $\mathcal{T}(N)$ is a tree-decomposition, $\mathrm{ori}(N)$ contains no $(\omega+1)$-chain. By [GK21, Corollary 3.3], this implies that $\mathrm{ori}(N)$ is the disjoint union of its splitting stars.

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 $v \in X(s_1) \cap X(s_2)$, we have $s_1 > s_2^*$, 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 $s_1 > s > s_2^*$. By definition of the partial ordering $\geqslant$, the vertex v lies 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

$$R_N(s) := R(s) \setminus \bigcup \{\, R(s') : s' \in \mathrm{ori}(N) \text{ such that } X(s') \text{ meets } X(s) \text{ and } s > s' \,\}$$

We remark that the clause of $X(s')$ meeting $X(s)$ in the definition of $R_N(s)$ is redundant, since any edge in $R(s) \cap R(s')$ has a vertex in $X(s) \cap X(s')$, as $s > s'$. We still include it as a reminder.

For the next lemmas, recall that the *interior* of a splitting star σ of N is defined as $\mathrm{int}(\sigma) := \bigcap_{(A,B)\in\sigma} B$. Whenever $\mathcal{T}(N)$ from Construction 2.2.3 is a tree-decomposition $(T, \mathcal{V})$ of G, its bags V_σ are $\mathrm{int}(\sigma)$ and we choose its parts G_σ as $G[V_\sigma]$.

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 $(T, \mathcal{V})$. Let σ be a splitting star of N and $s \in \sigma$. Then $R_N(s) = E(X(s), V_\sigma \setminus X(s))$.*

Proof. Let $E := E(X(s), V_\sigma \setminus X(s))$. For the backward inclusion $R_N(s) \supseteq E$, note first that $R(s) \supseteq E$ by the definition of V_σ. Consider $e = xy \in R(s) \setminus R_N(s)$ with $x \in X(s)$; we show that $e \notin E$. By the choice of e, there is $s' = (A, B) \in \mathrm{ori}(N)$ with $e \in R(s')$ such that $X(s)$ and $X(s')$ meet and $s > s'$; we choose s' maximal with respect to $>$. Then there is no $s'' \in \mathrm{ori}(N)$ with $s > s'' > s'$. By Lemma 2.2.5, we thus have $(s')^* \in \sigma$. But $x \in X(s')$, as $e \in R(s')$, and hence $y \in B \setminus A$. Thus, $y \notin V_\sigma$ and so $e \notin E$.

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 $\mathcal{K} := \mathcal{K}(N)$ the set of components of $G - \bigcup \mathcal{X}$ and define the relation $\Sigma_N^{\mathcal{K}}$ on $\mathrm{ori}(N)$ by letting

$$s_1 \Sigma_N^{\mathcal{K}} s_2 :\Leftrightarrow \text{there is a component } K \in \mathcal{K} \text{ such that } \partial K \text{ meets } R_N(s_i) \text{ for both } i = 1, 2.$$

Note that $\Sigma_N^{\mathcal{K}}$ is symmetric.

Lemma 2.8.3. *Let N be a nested set of proper separations of a connected graph G such that $\mathcal{T}(N)$ is a tree-decomposition $(T, \mathcal{V})$ with parts $G_\sigma = G[V_\sigma]$. Let $K \in \mathcal{K}(N)$. Then there is a unique node σ of T whose part G_σ contains K. Moreover, 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 are the separations in N, we can orient every edge of T in the direction of K to find a sink σ of T whose part G_σ contains K. Then σ is unique with this property.

Now let $s \in \mathrm{ori}(N)$ be given as stated. As ∂K meets $R_N(s) \subseteq R(s)$, our s points towards K. So we find, by the definition of splitting star, a unique $t \in \sigma$ with $s \geqslant t$. By assumption, $R_N(s)$ and ∂K share an edge xk with $x \in X(s)$ and $k \in K$. Since $X(t)$ is disjoint from K but $X(t)$ separates $X(s)$ from K, the edge xk is in $R(t)$. Hence, $s \not> t$ by $xk \in R_N(s)$, giving $s = t \in \sigma$ as desired. □

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 $s_1 \Sigma^{\mathcal{K}}_N s_2$. Then s_1 and s_2 lie in the same splitting star of N.* □

An edge e of G *bridges* a collection $\mathcal{X}$ of vertex sets of G if both ends of e are contained in $\bigcup \mathcal{X}$ but every $X \in \mathcal{X}$ contains at most one end of e. If N Is 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

$$s_1 \Sigma^{\mathcal{X}}_N s_2 :\Leftrightarrow \text{there is an } \mathcal{X}\text{-bridging edge of } G \text{ that lies in } R_N(s_i) \text{ for both } i = 1, 2.$$

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 $(T, \mathcal{V})$ with parts $G_\sigma = G[V_\sigma]$. Let e be an $\mathcal{X}(N)$-bridging edge of G. Then there is a unique node σ of T whose part G_σ contains e.*

Proof. As e bridges $\mathcal{X}(N)$ and the edges of T are the separations in N, we can orient every edge of T in the direction of e to find a sink σ of T whose part G_σ contains e. 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 the three symmetric relations $\Sigma^{\vec{x}_N}$, $\Sigma^{\mathcal{K}}_N$ 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 $(T, \mathcal{V})$ with parts $G_\sigma = G[V_\sigma]$. 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 the 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 $V_\sigma \setminus (X_1 \cup X_2)$, and thus avoid, by Lemma 2.2.5, all other $X \in \mathcal{X}$. So there is a (unique) component K of $G \setminus \bigcup \mathcal{X}$ that contains all the internal vertices of P. Let e_i denote the edge on P that is incident to x_i for both i; in particular, e_i is in ∂X_i and ∂K. Then $e_i \in R_N(s_i)$ by Lemma 2.8.2, so $K \subseteq G_\sigma$ by Lemma 2.8.3 and $s_i \in \sigma$ for $i = 1, 2$. Hence, the component K witnesses $s_1 \Sigma^{\mathcal{K}}_N s_2$, 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 $\mathcal{X}$-bridging and is an edge of the part G_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 $s_1 \Sigma^{\mathcal{X}}_N s_2$, 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 $s_1 \Sigma^{\vec{x}}_N s_2$ 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 *σ-part* of G, which we denote by $\mathrm{part}(\sigma)$ in formulas, is the subgraph of G that is the union of the graphs listed under (i) and (ii) below plus the edges listed under (iii):

(i) the induced subgraph $G[X(s)]$ for every $s \in \sigma$;

(ii) the components $K \in \mathcal{K}(N)$ 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 $s \in \sigma$.

Note that the edges listed in (iii) have both ends in graphs listed in (i) and (ii), by Lemma 2.8.2.

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 $(T, \mathcal{V})$ with parts $G_\sigma = G[V_\sigma]$. Let σ be a splitting star and, equivalently, a cutout of N. Then $G_\sigma = \mathrm{part}(\sigma)$.*

Proof. Let $I := G_\sigma$ and $P := \mathrm{part}(\sigma)$, and recall that $V_\sigma = \mathrm{int}(\sigma)$ for all nodes $\sigma \in T$ by the definition of $\mathcal{T}(N)$.

We first show the inclusion $P \subseteq I$. Since σ is a star, we have $X(s) \subseteq \mathrm{int}(\sigma)$ for every $s \in \sigma$, and in particular $G[X(s)] \subseteq I$. If K is a component in $\mathcal{K}(N)$ such that ∂K shares an edge with $R_N(s)$ for some $s \in \sigma$, then $K \subseteq G_\sigma = I$ by Lemma 2.8.3. The edges in $R_N(s)$ for any $s \in \sigma$ 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 $\mathcal{K}_\sigma$ denote the set of components of $I - \bigcup_{s \in \sigma} X(s)$. Fix any $K \in \mathcal{K}_\sigma$. By Lemma 2.2.5, $K \in \mathcal{K}(N)$ and, by Lemma 2.8.2, $\partial K \subseteq \bigcup_{s \in \sigma} R_N(s)$. Since G is connected, there is an edge $e \in \partial K$, so $K \subseteq P$ by (ii). It remains to show that every edge e of I that is not contained in a $G[X(s)]$ or ∂K for a $K \in \mathcal{K}_\sigma$ lies in P. By Lemma 2.2.5, such an edge e bridges $\mathcal{X}$. Moreover, e joins a vertex in $X(s_1)$ to a vertex in $X(s_2) \setminus X(s_1)$ for some two $s_1, s_2 \in \sigma$. By Lemma 2.8.2, e is in $R_N(s_1)$. Hence, $e \in P$ by (iii). □

2.8.2 Cutouts yield 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 let $s \succcurlyeq t$ 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 $s \succcurlyeq t$ if and only if $s^* \preccurlyeq t^*$.

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 $s_r \succcurlyeq t_r$.*

Proof. This follows from Correspondences 2.4.6 and 2.6.3 together with the 'moreover'-part of Lemma 2.2.2. □

Lift 2.8.10 * ($\geqslant$)**.** *Let G be a connected graph such that $(k-1)r < \Delta_r(G)$, and let s and t be oriented r-local separations of G with r-coupled r-tomic separators X and Y of size $\leqslant k$, respectively. Let $\hat{s}$ be any lift of s to G_r, with separator $\hat{X}$ say. Let $\hat{t}$ be the unique lift of t with separator $\hat{Y}$ such that $\hat{X}$ and $\hat{Y}$ are r-coupled, as considered in Lift 2.6.7*. Then $s \succ t$ if and only if $\hat{s} \succ \hat{t}$.*

Proof. The equivalence $s \succcurlyeq t \Leftrightarrow \hat{s} \succcurlyeq \hat{t}$ follows from Lift 2.6.7 *. Using the uniqueness of r-coupled lifts as in Lift 2.6.6 *, we derive the stronger equivalence $s \succ t \Leftrightarrow \hat{s} \succ \hat{t}$. □

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. For $s \in \mathrm{ori}(N)$ we write

$$R_N(s) := R(s) \setminus \bigcup \{\, R(s') : s' \in \mathrm{ori}(N) \text{ such that } X(s') \text{ meets } X(s) \text{ and } s \succ s' \,\}.$$

We remark that stating that $X(s')$ meets $X(s)$ is technically redundant, since it holds for any $s \succ s'$ such that $R(s) \cap R(s') \neq \emptyset$.

For the remainder of this section (2.8.2), let us 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 $N_r := \{s_r : s \in N\}$.

Correspondence 2.8.11. *For every $s \in \mathrm{ori}(N)$ we have $R_N(s) = R_{N_r}(s_r)$.*

Proof. This follows from Correspondences 2.4.6 and 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 $\mathcal{K} := \mathcal{K}(N)$ and define the relation $\Sigma_N^{\mathcal{K}}$ on $\mathrm{ori}(N)$ by letting

$$s_1 \Sigma_N^{\mathcal{K}} s_2 :\Leftrightarrow \text{there is a component } K \in \mathcal{K} \text{ such that } \partial K \text{ meets } R_N(s_i) \text{ for both } i = 1, 2.$$

Note that $\Sigma_N^{\mathcal{K}}$ is symmetric. We define the relation $\Sigma_N^{\mathcal{X}}$ on $\mathrm{ori}(N)$ by letting

$$s_1 \Sigma_N^{\mathcal{X}} s_2 :\Leftrightarrow \text{there is an } \mathcal{X}\text{-bridging edge of } G \text{ that lies in } R_N(s_i) \text{ for both } i = 1, 2.$$

Note that $\Sigma_N^{\mathcal{X}}$ is symmetric. As for ordinary separations, we denote by Σ_N the transitive reflexive closure of the union of all three symmetric relations $\Sigma^{\vec{x}_N}$, $\Sigma_N^{\mathcal{K}}$ and $\Sigma_N^{\mathcal{X}}$. 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 *σ-part* of G for a cutout σ of N is defined just like for a cutout of a set of usual separations before Proposition 2.8.8, 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 $* \in \{\vec{x}, \mathcal{K}, \mathcal{X}\}$.*

(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 σ-parts of G coincide with its σ_r-parts.*

Proof. Note that $\mathcal{X}(N) = \mathcal{X}(N_r)$ and $\mathcal{K}(N) = \mathcal{K}(N_r)$.

(i). For $* = \vec{x}$ this is immediate from Correspondence 2.8.9. For $* \in \{\mathcal{X}, \mathcal{K}\}$ it follows from Correspondence 2.8.11.

(ii) follows from (i).

(iii) follows from (ii) by Correspondence 2.8.11. □

2.8.3 Interplay with the covering map

Projection 2.8.13 *. *Let G be a connected graph such that $(k-1)r < \Delta_r(G)$, and let N be a nested set of r-local separations of G with r-tomic separators of size $\leqslant k$. Let $\hat{N}$ be the set of all lifts of elements of N from Lift 2.4.19 and Lift 2.5.10*. Then, for every $\hat{s} \in \mathrm{ori}(\hat{N})$ and $s := p_r(\hat{s})$, the r-local covering p_r restricts to a bijection $R_{\hat{N}}(\hat{s}) \to R_N(s)$.*

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 map p_r induces a bijection $\hat{S} \to S$, and for every $\hat{t} \in \hat{S} \cup \{\hat{s}\}$ 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 G be a connected graph such that $(k-1)r < \Delta_r(G)$, let N be a nested set of r-local separations of G with r-tomic separators of size $\leqslant k$, and let $\hat{N}$ be the set of all lifts of elements of N.*

(i) *For all* $\hat{s},\hat{t} \in \mathrm{ori}(\hat{N})$ *and* $* \in \{\vec{x},\mathcal{K},\mathcal{X}\}$ *we have* $\hat{s}\Sigma^*_{\hat{N}}\hat{t} \Longrightarrow p_r(\hat{s})\Sigma^*_N\, p_r(\hat{t})$.

(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 $\hat{s},\hat{t} \in \mathrm{ori}(\hat{N})$ be 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 2.8.10 *. For $* = \mathcal{X}$ the assertion follows from Projection 2.8.13 * and $\bigcup\mathcal{X}(\hat{N}) = p_r^{-1}(\bigcup\mathcal{X}(N))$.

For $* = \mathcal{K}$ there is a component $\hat{K}$ of $G_r - \bigcup\mathcal{X}(\hat{N})$ such that $\partial\hat{K}$ meets $R_{\hat{N}}(\hat{s})$ and $R_{\hat{N}}(\hat{t})$. By path lifting, the projection $K := p_r(\hat{K})$ is an entire component of $G - \bigcup\mathcal{X}(N)$. Then both $R_N(s)$ and $R_N(t)$ meet ∂K by Projection 2.8.13 *, giving $s\Sigma^{\mathcal{K}}_N t$.

(ii). The existence of σ given $\hat{\sigma}$ is immediate from (i). Every component $\hat{K} \in \mathcal{K}(\hat{N})$ projects to a component $p_r(\hat{K}) \in \mathcal{K}(N)$, as observed in the proof of (i). Hence, the 'moreover'-part follows with Projection 2.8.13 *. □

Lift 2.8.15 * (Σ^*_N). *Let* G *be a connected graph such that* $(k-1)r < \Delta_r(G)$, *let* N *be a nested set of* r*-local separations of* G *with* r*-tomic separators of size* $\leqslant k$, *and let* $\hat{N}$ *be the set of all lifts of elements of* N.

(i) *Let* $s,t \in \mathrm{ori}(N)$ *and* $* \in \{\vec{x},\mathcal{K},\mathcal{X}\}$ *be 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,* $\mathrm{part}(\sigma) \subseteq p_r(\mathrm{part}(\hat{\sigma}))$.

We remark that the lift $\hat{t}$ of t in (i) will not normally be unique when $* = \mathcal{K}$.

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 assume that $X(s)$ and $X(t)$ meet in a vertex v, and we consider the unique lift $\hat{t}$ of t such that $X(\hat{t})$ meets $X(\hat{s})$ (in a lift $\hat{v}$ of v) from the statement of Lift 2.6.7 *. Then $\hat{s} \succ \hat{t}^*$ by Lift 2.8.10 *. If there is $\hat{s}' \in \mathrm{ori}(\hat{N})$, which is a lift of some $s' \in \mathrm{ori}(N)$, with $\hat{v} \in X(\hat{s}')$ and $\hat{s} \succ \hat{s}' \succ \hat{t}^*$, then $v \in X(s')$ and Lift 2.8.10 * yields $s \succ s' \succ t^*$, contradicting $s\Sigma^{\vec{x}}_N t$. Hence we have $\hat{s}\Sigma^{\vec{x}}_{\hat{N}}\hat{t}$.

For $* = \mathcal{K}$ consider any $K \in \mathcal{K}(N)$ 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 $\hat{s}\Sigma^{\mathcal{K}}_{\hat{N}}\hat{t}$.

For $* = \mathcal{X}$ there is an $\mathcal{X}(N)$-bridging edge e of G in $R_N(s) \cap R_N(t)$. Its lift $\hat{e}$ at the given $\hat{s}$ bridges $\mathcal{X}(\hat{N})$, because e bridges $\mathcal{X}(N)$, and it lies in $R_{\hat{N}}(\hat{s}) \cap R_{\hat{N}}(\hat{t})$ 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 $*(i) \in \{\vec{x},\mathcal{K},\mathcal{X}\}$ with $s_{i-1}\Sigma^{*(i)}_N s_i$, where $s_0 := s$ and $s_n := t$. Let $\hat{s}_0 := \hat{s}$. Applying (i) iteratively for $i = 1,\dots,n$ yields lifts $\hat{s}_i$ of s_i such that $\hat{s}_{i-1}\Sigma^{*(i)}_N\hat{s}_i$. Thus, $\hat{t} = \hat{s}_n \in \hat{\sigma}$, as desired. The 'moreover'-part follows with Projection 2.8.13 *, since components in $\mathcal{K}(\hat{N})$ project to components in $\mathcal{K}(N)$. □

Proposition 2.8.16 * (Cutout). *Let G be a connected graph such that $(k-1)r < \Delta_r(G)$, let N be a nested set of r-local separations of G with r-tomic separators of size $\leqslant k$, and let $\hat{N}$ be the set of all lifts of elements of 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, $p_r(\text{part}(\hat{\sigma})) = \text{part}(p_r(\hat{\sigma}))$.*

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). The converse inclusion follows from Lift 2.8.15 * (ii), so $\sigma = p_r(\hat{\sigma})$.

Surjective. Let σ be any 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). As we just saw, $p_r(\hat{\sigma})$ is itself a cutout of N, which must be σ since the cutouts of N partition Σ_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 with parts, as follows. If $N = \emptyset$, let $(H, \mathcal{G})$ with $\mathcal{G} = (G_h : h \in V(H))$ 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 decomposition part G_h for a node h of H to be the h-part of G, with h viewed as a cutout of N.

Note that $\mathcal{H}(N)$ is uniquely determined by G and N. Thus, if N is Γ-canonical for some subgroup Γ of $\text{Aut}(G)$, then so is $\mathcal{H}(N)$. Additional assumptions on N are required to ensure that $\mathcal{H}(N)$ is a graph-decomposition.

Recall that an isomorphism $\psi\colon G \to G'$ between (multi-)graphs 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.

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 $\mathcal{T}(N)$ is a tree-decomposition $(T, \mathcal{V})$ with parts $P_t = G[V_t]$. Let $(H, \mathcal{G}) := \mathcal{H}(N_r)$ with parts $G_h = \text{part}(h)$. Then the map $s \mapsto s_r$ induces a canonical isomorphism $\psi\colon T \to H$ such that $G_{\psi(\sigma)} = P_\sigma$ 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 combines with the $E(T) \to E(H)$ map $s \mapsto s_r$ to a graph isomorphism $\psi\colon T \to H$. We have $\text{part}(\sigma) = P_\sigma$ by Proposition 2.8.8, and we have $\text{part}(\sigma) = \text{part}(\sigma_r)$ by Correspondence 2.8.12 (iii). Furthermore, $\text{part}(\sigma_r) = G_{\sigma_r}$ by Construction 2.8.17. With $\psi(\sigma) = \sigma_r$ we find that $P_\sigma = G_{\psi(\sigma)}$. □

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, 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, there are only finitely many separations in N whose separators contain v.*

(ii) *$\mathcal{T}(N)$ is a tree-decomposition of G.*

Proof. (i). Let N_v consist of all $s \in N$ with $v \in X(s)$. Let $U := B_G(v, (k-1)\lfloor r/2 \rfloor + 1)$. 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 every $s \in N_v$. 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 [EKT22b, Lemma 2.7] it suffices to show that N has 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 prove this, use (i) to find an infinite subsequence with disjoint separators as in the proof of [DJKK26, Lemma 7.5]. □

Lemma 2.8.20. *Let G be a locally finite, connected graph, and assume that the binary cycle space of G is generated by cycles of length $\leqslant r$. Let N be a nested set of r-local separations of G with r-tomic separators all of size $\leqslant k$ for some $k \in \mathbb{N}$. Then $\mathcal{H}(N)$ is a tree-decomposition (with parts) of G. Moreover, if N is Γ-canonical for some subgroup Γ of $\mathrm{Aut}(G)$, then $\mathcal{H}(N)$ is also Γ-canonical.*

Proof. By Correspondence 2.4.6, there exists a unique set M of separations of G with r-tomic separators such that $M_r = N$. Since all separations in M have r-tomic separators of size $\leqslant k$, we have that $\mathcal{T}(M)$ is a tree-decomposition by Lemma 2.8.19 (ii). Thus, $\mathcal{H}(N)$ is as desired by Correspondence 2.8.18. □

Theorem 2.8.21 *. *Let G be a locally finite, connected graph such that $(k-1)r < \Delta_r(G)$. Let N be a nested set of r-local separations of G with r-tomic separators of size $\leqslant k$, and let $(H, \mathcal{G}) := \mathcal{H}(N)$, with $\mathcal{G} = (G_h : h \in V(H))$ say. Let $\hat{N}$ be the set of all lifts of all elements of N to G_r, and recall that $\mathcal{H}(\hat{N}) =: (\hat{H}, \hat{\mathcal{G}})$ is a Γ_r-canonical tree-decomposition (with parts) of G_r by Lemma 2.8.20. Let $\bar{\mathcal{H}}(N) := (\bar{H}, \bar{\mathcal{G}})$ with family $\bar{\mathcal{G}} = (\bar{G}_h : h \in V(\bar{H}))$ of parts be the graph-decomposition of G that is obtained from $\mathcal{H}(\hat{N})$ by folding 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 $\tau(s)$ denote the set of all lifts of s. Then τ witnesses that $\mathcal{H}(N)$ and $\bar{\mathcal{H}}(N)$ are isomorphic, that is,

(i) *the pair of bijections $\tau := (\sigma \mapsto \tau(\sigma), s \mapsto \tau(s))$ is an isomorphism $H \to \bar{H}$, and*

(ii) *$G_\sigma = \bar{G}_{\tau(\sigma)}$ for every $\sigma \in V(H)$.*

In the context of Theorem 2.8.21 *, we recall that the lifts of the $s \in N$ exist by Lift 2.5.10 * and Lemma 2.6.5 and have r-tomic separators by Lemma 2.5.7 *, and $\mathcal{H}(\hat{N})$ is a tree-decomposition (with parts) of G_r by Lemma 2.8.20.

Proof. (i). We first show that the map $\sigma \mapsto \tau(\sigma)$ is a bijection $V(H) \to V(\bar{H})$. The vertices of $\bar{H}$ are the orbits of the vertices of $\hat{H}$ under the action of the group $\Gamma_r := \Gamma_r(G)$ of deck transformations of p_r on $V(\hat{H})$. The vertices of $\hat{H}$, in turn, are the cutouts of $\hat{N}$. The group Γ_r acts naturally on $\hat{N}$, and since deck transformations preserve Σ_N, Γ_r also acts naturally on the cutouts of $\hat{N}$. By Proposition 2.8.16* p_r maps the cutouts of $\hat{N}$ onto the cutouts of N, and this map is well-defined on the Γ_r-orbits of cutouts of $\hat{N}$. By Lift 2.5.10* the set of all lifts of a separation in N forms a Γ_r-orbit of $\hat{N}$, and by Lift 2.8.15* (ii) this transfers to cutouts of N and their lifts to $\hat{N}$. Thus, the Γ_r-orbits of the set of cutouts of $\hat{N}$ are the sets $\tau(\sigma)$ for cutouts σ of N. All in all, p_r maps the sets $\tau(\sigma)$, the vertices of $\bar{H}$, bijectively to the cutouts of N, the vertices of H.

Similarly, the edges of H are the separations in N, while the edges of $\bar{H}$ are the Γ_r-orbits of the edges of $\hat{H}$, the separations in $\hat{N}$. The group Γ_r acts naturally on $\hat{N}$, and by Projection 2.5.9* p_r maps the Γ_r-orbits of $\hat{N}$ to N. As above the Γ_r-orbits of $\hat{N}$ are the sets $\tau(s)$ for separations s in N. Thus, p_r maps the sets $\tau(s)$, the edges of $\bar{H}$, bijectively to the separations in N, the edges of H.

The incidences between edges and vertices are clearly preserved.

(ii). Let $\sigma \in V(H)$ be a cutout. We have to show $G_\sigma = \bar{G}_{\tau(\sigma)}$. We have already seen in (i) that $\tau(\sigma)$ is a Γ_r-orbit of the cutouts of $\hat{N}$. As $\bar{\mathcal{H}}(N)$ is obtained from the tree-decomposition $\mathcal{H}(\hat{N})$ (with parts that are induced subgraphs of G_r) by folding via p_r as in Construction 2.2.6, we therefore have for any cutout $\hat{\sigma} \in \tau(\sigma)$ of $\hat{N}$ that $\bar{G}_{\tau(\sigma)} = p_r(\hat{G}_{\hat{\sigma}}) = p_r(\text{part}(\hat{\sigma}))$. Finally, $p_r(\text{part}(\hat{\sigma})) = \text{part}(\sigma) = G_\sigma$ by Proposition 2.8.16* using $\sigma = p_r(\hat{\sigma})$. □

Corollary 2.8.22*. *Let G be a locally finite, connected graph such that $(k-1)r < \Delta_r(G)$. Let N be a nested set of r-local separations of G with r-tomic separators of size $\leqslant k$. Then $\mathcal{H}(N)$ is a graph-decomposition with parts. Moreover, if N is Γ-canonical for a subgroup Γ of $\mathrm{Aut}(G)$, then so is $\mathcal{H}(N)$.* □

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 $\tilde{G}_{\tilde{h}}$ of G 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, since the edge set of H is, formally, the set N, of which N' is a subset.

Lemma 2.8.23*. *Let G be a locally finite, connected graph such that $(k-1)r < \Delta_r(G)$. Let $N \subseteq M$ be two nested sets of r-local separations of G with r-tomic separators of size $\leqslant k$. Then $\mathcal{H}(N) = \mathcal{H}(M)/(M \smallsetminus N)$.*

Proof. Let $\hat{N}$ be the set of all lifts to G_r of the elements of N, which is possible by Lift 2.5.10* 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]. Since $\hat{N}$ is a Γ_r-invariant subset of $\hat{M}$, this construction is invariant under Γ_r itself and hence p_r is defined on Γ_r-orbits, to yield by Construction 2.2.6 graph-decompositions $\bar{\mathcal{H}}(N)$ and

$\bar{\mathcal{H}}(M)$ of G such that $\bar{\mathcal{H}}(N) = \bar{\mathcal{H}}(M)/(M \smallsetminus N)$. By Theorem 2.8.21 *, we then get $\mathcal{H}(N) = \bar{\mathcal{H}}(N) = \bar{\mathcal{H}}(M)/(M \smallsetminus N) = \mathcal{H}(M)/(M \smallsetminus N)$. □

2.9 Proofs of Theorem 6 and Theorem 7

Recall that $N_r^{\leqslant k}(G)$ was defined in Construction 2.7.16, which for this section we take to refer to the set of *all* r-local bottlenecks of order $\leqslant k$ in G as the set $\mathcal{B}_r^k(G)$. We define $\mathcal{H}_r^{\leqslant k}(G) := \mathcal{H}(N_r^{\leqslant k}(G))$ with the operator $\mathcal{H}$ taken 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 6. *Let G be a finite connected graph and $r, k \in \mathbb{N}$ with $k < K(G,r) = \frac{\Delta_r(G)}{r} + 1$.*

(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 canonical tangle tree-decomposition of scope 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 canonical 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. By assumption we have $(k-1)r < \Delta_r(G)$. Further, we claim that we may assume without loss of generality that * is satisfied; that is, $r \geqslant 2$ or $k = 1$. 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 $k = 1$.

(ii). By Correspondence 2.7.17, $N_r^{\leqslant k}(G_r)$ is well-defined and nested, and by Construction 2.7.16, the r-local separations in $N_r^{\leqslant k}(G_r)$ were selected from r-local bottlenecks. So they are tight and, in particular, r-tomic by Lemma 2.4.2. Hence, $\mathcal{H}_r^{\leqslant k}(G_r)$ is a tree-decomposition of G_r by Lemma 2.8.20, and it is canonical because $N_r^{\leqslant k}(G_r)$ is canonical by Correspondence 2.7.17. 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, $N_r^{\leqslant k}(G_r)$ is a tree of tangles of G_r of scope k, and so $\mathcal{H}_r^{\leqslant k}(G_r)$ is a tangle tree-decomposition of scope k of G_r.

(iii). By Theorem 2.7.18 *, $N_r^{\leqslant k}(G)$ is well-defined and nested, and $N_r^{\leqslant k}(G_r)$ equals the set of all lifts to G_r of all elements of $N_r^{\leqslant k}(G)$. Further, the elements of $N_r^{\leqslant k}(G)$ are r-tomic by Lemma 2.4.2 because they come from r-local bottlenecks by Construction 2.7.16. Hence, (iii) follows from (ii) and Theorem 2.8.21 *.

(i). By (iii) $\mathcal{H}_r^{\leqslant k}(G)$ is a graph-decomposition of G. Since $N_r^{\leqslant k}(G)$ is canonical by Theorem 2.7.18 * and $\mathcal{H}_r^{\leqslant k}(G) = \mathcal{H}(N_r^{\leqslant k}(G))$, also $\mathcal{H}_r^{\leqslant k}(G)$ is canonical by Construction 2.8.17.

So it remains to show that the edges of its decomposition graph H represent all the r-local bottlenecks in G of order $\leqslant k$. Recall that $E(H) = N_r^{\leqslant k}(G)$ by Construction 2.8.17. And indeed, every r-local bottleneck β_r in G of order $\leqslant k$ shares an r-local separation with $N_r^{\leqslant k}(G)$ 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 7. Given G, r, k as in Theorem 6, 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 $\mathcal{H}_r^{\leqslant k}(G)$. □

2.A Corners of local separations

We did not require an analogue of corners for local separations in this chapter. Corners of genuine separations are frequently used in the literature, however, and we believe that their local counterparts can be defined as follows. The following definitions are supported by Figure 2.18. Let $\{A_1, A_2\}$ and $\{C_1, C_2\}$ be two separations of a graph G with separators X and Y, respectively. Let $i, j \in [2]$. Recall that the *corner-separator for A_i and C_j* is the union of the Y-link for A_i with the X-link for C_j with $X \cap Y$. And the *corner for A_i and C_j* is the separation $\{A_i \cap C_j, A_{3-i} \cup C_{3-j}\}$. As i, j range over $[2]$, there are four corner-separators and four corners. Now let $\{E_1, X, E_2\}$ and $\{F_1, Y, F_2\}$ be two oriented r-local separations of G. The *corner-separator for E_i and F_j* is the union of the Y-link for E_i with the X-link for F_j with $X \cap Y$. Let L denote the corner-separator for E_i and F_j. The *corner for E_i and F_j* is the triple

$$\{\, \partial L \smallsetminus (E_{3-i} \cup F_{3-j}),\ L,\ \partial L \cap (E_{3-i} \cup F_{3-j}) \,\}.$$

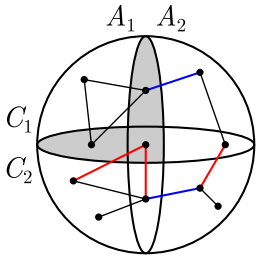

Figure 2.18: The edges in E_2 are coloured blue, while the edges in F_2 are coloured red. The corner-separator L is coloured grey

Correspondence 2.A.1 (Corners)**.** *Assume that the binary cycle space of G is generated by cycles of length $\leqslant r$ and that G is connected. Let $s = \{A_1, A_2\}$ and $s' = \{C_1, C_2\}$ be crossing separations of G with r-tomic separators. Let $\{E_1, X, E_2\} := s_r$ and $\{F_1, Y, F_2\} := s'_r$, as well as $i, j \in [2]$. The corner d of s_r, s'_r for E_i and F_j is equal to c_r where c is the corner of s, s' for A_i and C_j. In particular, the four corners of s_r, s'_r are r-local separations of G.*

Proof. Without loss of generality $i, j = 1$. Then $c = \{A_1 \cap C_1, A_2 \cup C_2\}$. The separator L of c is equal to the separator of d by Correspondence 2.6.3. It therefore remains to show

the following two equalities:

$$\partial L \smallsetminus (E_2 \cup F_2) = E(L, (A_1 \cap C_2) \smallsetminus L)$$
$$\partial L \cap (E_2 \cup F_2) = E(L, (A_2 \cup C_2) \smallsetminus L)$$

These are straightforward to check, for example by using Figure 2.18. □

CHAPTER 3

Canonical tree-decompositions of chordal graphs

This chapter is based on [JK25] and joint work with Paul Knappe.

Halin characterised the chordal locally finite graphs as those that admit a tree-decomposition into cliques. We show that these tree-decompositions can be chosen to be canonical, that is, so that they are invariant under all the graph's automorphisms. As an application, we show that a locally finite, connected graph G is r-locally chordal (that is, its $r/2$-balls are chordal) if and only if the unique canonical graph-decomposition $\mathcal{H}_r(G)$ of G which displays its r-global structure is into cliques. Our results also serve as tools for further characterisations of r-locally chordal graphs.

3.1 Introduction

It is well-known[1] that graphs admitting tree-decompositions into cliques[2] are chordal. Here, we say that a tree-decomposition $(T, \mathcal{V})$ of a graph G is *into cliques* if all its bags V_t are cliques in G, that is, every $G[V_t]$ is a complete subgraph of G.

Proposition 3.1.1 (Folklore)**.** *Every (possibly infinite) graph that admits a tree-decomposition into cliques is chordal.*

For finite graphs, it is also well-known that the converse of Proposition 3.1.1 holds, too (see e.g. Gavril [G74, Theorem 2 and 3]): the chordal graphs are precisely those graphs admitting tree-decompositions into cliques. This result was later extended to graphs with no infinite clique, such as locally finite graphs, by Halin [H64, 10.2'] based on his *zweiter Zerlegungssatz* [H89, Satz 5]:

Theorem 3.1.2 (Halin)**.** *A graph G with no infinite clique is chordal if and only if it admits a tree-decomposition into (its maximal) cliques.*

Diestel [D88, Theorem 4.2] even characterised the countable graphs that admit a tree-decomposition into cliques via forbidden 'simplicial minors'. This characterisation does not extend to uncountable graphs [D91, Paragraphs after Theorem 1.4], but Diestel conjectured an extension of his characterisation [D91, End of §1].

[1] This statement is immediate; e.g. see the first paragraph of the proof presented in [G74, Theorem 3].

[2] In this chapter, we will ask whether a graph admits certain tree-decompositions into cliques. We remark that all these questions can equivalently be formulated in the language of subtree intersection graphs (over trees); see [AK25].

When working with highly symmetric graphs, such as Cayley graphs or covers of (finite) graphs, it is often useful to find so-called *canonical* tree-decompositions that are invariant under the automorphism-group of the graph (see e.g. [CHM22; DJKK26; DK15; HLMR22]). Formally, a tree-decomposition $(T, \mathcal{V})$ of a graph G is Γ-*canonical* for a subgroup Γ of the automorphism-group $\mathrm{Aut}(G)$ of G if there is an action of Γ on T that commutes with $t \mapsto V_t$, i.e. $V_{\varphi \cdot t} = \varphi(V_t)$ for every $t \in V(T)$ and $\varphi \in \Gamma$; usually, the action of Γ on T is unique (Lemma 3.2.7). If $(T, \mathcal{V})$ is $\mathrm{Aut}(G)$-canonical, then we call it *canonical* for short.

Does Theorem 3.1.2 still hold if we want the tree-decomposition into cliques to be canonical? In other words, does every chordal graph with no infinite clique admit a canonical tree-decomposition into cliques? Our main result answers this question in the affirmative for chordal graphs that are locally finite and connected[3]:

Theorem 8. *Let G be a locally finite, connected graph. Then G is chordal if and only if G admits a canonical tree-decomposition into cliques.*

Theorem 3.1.2 makes the even stronger assertion that there are tree-decompositions *into the graph's maximal cliques*. Formally, a *maximal* clique of a graph G is an inclusion-wise maximal clique, and a tree-decomposition $(T, \mathcal{V})$ of G is *into its maximal cliques* if the map $t \mapsto V_t$ is a bijection from $V(T)$ to the set of maximal cliques of G. We show that Theorem 8 cannot be strengthened to canonical tree-decompositions into maximal cliques: there are chordal graphs, even finite and connected ones, that do not admit a canonical tree-decomposition into their maximal cliques (Example 3.5.1).

In contrast, we provide a positive result for graphs $\hat{G}$ from normal coverings $p : \hat{G} \to G$ for locally finite, connected graphs G (see Section 3.3.2 for the definitions). If such $\hat{G}$ is chordal, then it admits a tree-decomposition into its maximal cliques which is $\Gamma(p)$-canonical for the group $\Gamma(p)$ of deck transformations of p:

Theorem 9. *Let G be a locally finite, connected 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.*

We finally discuss what happens behind the barrier of locally finite graphs: there are countable chordal graphs with no infinite clique that admit a tree-decomposition into cliques, but no canonical one (Example 3.5.2), and there are countable chordal graphs that admit a tree-decomposition into cliques, but no tree-decomposition into their maximal cliques (Example 3.5.3).

3.1.1 Application: Characterizing locally chordal graphs by canonical graph-decompositions via coverings

A recent development in structural graph theory aims to study the local structure of graphs through the lens of local coverings. Introduced by Diestel, Jacobs, Knappe and Kurkofka in [DJKK26], the *r-local covering* $p_r : G_r \to G$ of a graph G is a (topological) covering

[3] Disconnected chordal graphs sometimes, e.g. the disjoint union of two paths of length 3, do not admit *canonical* tree-decompositions into cliques for the simple reason that tree-decompositions do not allow forests as model graphs. We remark that a disconnected chordal graph admits a canonical forest-decomposition into cliques if and only if its components admit canonical tree-decompositions into cliques. Hence, we restrict our view in this chapter to connected graphs and tree-decompositions.

of G that maintains all r-local structure of G while unfolding every other structure in a tree-like way (see Section 3.4.1 for the formal definition); here, the integer parameter $r \geqslant 0$ describes the intended degree of locality.

The main result of [DJKK26] applies state-of-the-art tangle-theory to the r-local cover G_r of G to obtain a canonical tree-decomposition $\mathcal{T}(G_r)$ of G_r which can then be re-folded along the covering map $p_r : G_r \to G$ to obtain a canonical graph-decomposition $\mathcal{H}_r(G)$ of G that displays the r-global structure of G (see Sections 3.4.2 and 3.4.4).

Theorem 3.1.3. ([DJKK26, Theorem 1]) *Let G be a finite graph, and $r > 0$ an integer. Then G has a unique canonical decomposition $\mathcal{H}_r(G)$ modelled on another finite graph $H = H(G, r)$ that displays its r-global structure.*

A different construction of graph-decompositions like $\mathcal{H}_r(G)$ is given by Jacobs, Knappe and Kurkofka in [JKK25]. It avoids the use of the usually infinite r-local covering and instead builds on a theory of r-local separators.

One of the main questions arising from Theorem 3.1.3 is whether one can characterize classes of graphs G by their decompositions $\mathcal{H}_r(G)$. We answer this question in the affirmative for the class of 'r-locally chordal' graphs, hereby providing the first known such characterisation. Given an integer $r \geqslant 0$, a graph is *r-locally chordal* if all its balls of radius $r/2$ are chordal (cf. Section 3.4.1).

Theorem 10. *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 $\mathcal{H}_r(G)$ of G which displays its r-global structure is into cliques.*

For the proof of Theorem 10, we build on [AKK25, Theorem 1] which, among other characterisations, asserts that the r-locally chordal graphs G are precisely those whose r-local cover G_r is chordal. We can then use Theorem 8, or more precisely its proof, to show that for connected, locally finite, chordal graphs such as G_r, the canonical tree-decomposition $\mathcal{T}(G_r)$ is into cliques. This transfers to the graph-decomposition $\mathcal{H}_r(G)$ obtained from $\mathcal{T}(G_r)$ by folding along p_r (Lemma 3.4.3). In a similar way, the even stronger Theorem 9 can be used to obtain canonical graph-decompositions of an r-locally chordal graph G which are into maximal cliques and even 'r-acyclic'; see Section 3.4.5.

3.1.2 Structure of the chapter

In Section 3.2 we recall the definitions and concepts around separators, separations and tree-decompositions and prove a few simple lemmas for later use. We then prove our main results, Theorems 8 and 9, in Section 3.3. In Section 3.4, we discuss the application of our results to r-locally chordal graphs and prove Theorem 10. Finally, we provide counterexamples to various strengthenings of Theorem 8 in Section 3.5.

Acknowledgements

We thank Tara Abrishami for many stimulating discussions on locally chordal graphs and local coverings. We thank Mai Trinh for useful comments on the chapter.

3.2 Background

We follow standard graph theory notation and definitions from [D17b]. We also refer to [H02] for terminology on topology.

In this chapter we deal with simple graphs, i.e. graphs with neither loops nor parallel edges. Often, we will restrict our view to locally finite graphs; a graph is *locally finite* if each vertex has finite degree. We specify when we deal with finite, locally finite, or general (possibly infinite) graphs.

For a set $\mathcal{X}$ of sets, we denote $\bigcup_{X\in\mathcal{X}} X$ by $\bigcup\mathcal{X}$.

3.2.1 Minimal, tight and efficient separators

Let u, v be two vertices and X a vertex-subset of a graph G. Then we say that X *separates* u *and* v in G and call X an *u–v separator* of G if $X \subseteq V(G) \setminus \{u, v\}$ and every u–v path in G meets X. If X separates some two vertices in G, then X is a *separator* of G. A separator X is a *minimal (vertex) separator* of G if there are two vertices u, v of G such that X is inclusion-minimal among the u–v separators in G.

Dirac characterized chordal graphs by their minimal separators:

Theorem 3.2.1 (Dirac [D61], Theorem 1)**.** *A (possibly infinite)*[4] *graph is chordal if and only if every minimal separator is a clique.*

A vertex-subset X of a graph G is a *tight* separator of G if at least two components C of $G - X$ are *full*, i.e. $N_G(C) = X$.[5] It is easy to see that a separator X of G is tight if and only if it is minimal. The relation between a minimal (equivalently: tight) finite separator in a chordal graph and the respective full components is quite restrictive; indeed, we can extract the following structure from a proof of Hofer-Temmel and Lehner [HL18]:

Lemma 3.2.2. *Let X be a finite clique in a (possibly infinite) chordal graph G. Then for every full component C of $G - X$, there is a vertex in C complete to X in G.*

Proof. Follow Hofer-Temmel and Lehner's proof of [HL18, Lemma 14] to obtain the desired result; there, the existence of the paths P_w is guaranteed for every $w \in X$ by C being a full component of $G - X$. □

We remark that Lemma 3.2.2 fails if we allow the clique X to be infinite[6]: Consider the chordal graph G consisting of two countable infinite cliques $X = \{x_0, x_1, \dots\}$ and $C = \{c_0, c_1, \dots\}$ that are joined by all edges $x_i c_j$ with $i \leqslant j$. Then C is a full component

[4] It is immediate from the proof presented in [BP93, Theorem 2.1] that this also holds for infinite graphs.

[5] Full components are sometimes also called *tight* components.

[6] The statement of [HL18, Lemma 3.13] does not include this necessary assumption of X being finite. If X is infinite, then their proof fails in that there might not even exist a single inclusion-maximal element among the C_v for $v \in V_1$.

of $G - X$, but every $c_j \in X$ has only the finite set $\{x_i : i \leqslant j\}$ of neighbours in X, contradicting the statement.

Finally, two maximal cliques in a chordal graph can also be told apart by vertex sets of comparatively smaller size. To formalize this, let A, B, X be vertex-subsets of a graph G. Then we say that X *separates* A *and* B in G and call X an *A–B separator* of G if each A–B path in G meets X. If the size of X is minimal among all the A–B separators, then X *efficiently* separates A and B.

Now an efficient separator of any two finite, maximal cliques in a chordal graph is smaller than the minimum of their sizes:

Theorem 3.2.3. *Let X and Y be two distinct maximal cliques in a (possibly infinite) chordal graph G. If at least one of X and Y is finite, then there is a X–Y separator of G whose size is less than the sizes of X and Y.*

As a tool for the proof of Theorem 3.2.3, we need the following version of Menger's Theorem for infinite graphs, which is due to Erdős and was published in [K36].

Theorem 3.2.4. *Let A and B be vertex-subsets of a (possibly infinite) graph G. Then the cardinality of an efficient A–B separator of G (equivalently: the minimum cardinality of an A–B separator) is the maximum cardinality of a family $\mathcal{P}$ of disjoint A–B paths in G.*

Proof of Theorem 3.2.3. By symmetry of the statement in X and Y, we may assume without loss of generality that $|X| \leqslant |Y|$. Since X and Y are distinct maximal cliques, there is a vertex $y \in Y \setminus X$. Thus, the component C of $G - X$ which meets Y is non-empty; note that C is unique and contains $Y \setminus X$, since Y is a clique.

Suppose for a contradiction that there is no X–Y separator of size less than $|X|$. We claim that C then is a full component of $G - X$. Indeed, by Theorem 3.2.4, there is a family of $|X|$ disjoint X–Y paths in G, and all of them must be contained in C except for their endvertices in X, thus witnessing that C is a full component of $G - X$. Now we apply Lemma 3.2.2 to obtain a vertex $w \in C$ complete to X. This then contradicts the maximality of the clique X in G, as $X \cup \{w\}$ forms a larger clique. □

Corollary 3.2.5. *Every efficient separator of two distinct maximal cliques in a (possibly infinite) chordal graph is minimal and tight.*

Proof. Let S be an efficient separator of two maximal cliques X and Y. By Theorem 3.2.3, S has size less than the sizes of X and Y. Thus, there exist vertices $x \in X \setminus S$ and $y \in Y \setminus S$. Then S is also a x–y separator, and it is in fact a minimal one since each x–y separator also separates $X \ni x$ and $Y \ni y$ and S did so efficiently. □

3.2.2 Tree-decompositions

Let G be a graph, let T be a tree, and let $\mathcal{V} = (V_t \mid t \in V(T))$ be a family of vertex-subsets of G indexed by the nodes of T. The pair $(T, \mathcal{V})$ is a *tree-decomposition* of G if

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

(T2) for every vertex v of G, the subgraph T_v of T induced by the set of nodes of t whose bags V_t contain v is connected.

For a tree-decomposition $(T, \mathcal{V})$, we call T its *decomposition tree*, the V_t its *bags* and the induced subgraphs $G[V_t]$ its *parts*. For an edge $f = t_1t_2 \in E(T)$, the set $V_f \coloneqq V_{t_1} \cap V_{t_2}$ is the *adhesion set* corresponding to f. We often introduce a tree-decomposition as a pair $(T, \mathcal{V})$ and tacitly assume that $\mathcal{V} = (V_t \mid t \in V(T))$.

Roughly speaking, (T1) says that $\mathcal{V}$ forms a decomposition of G, and, by (T2), the decomposition $\mathcal{V}$ follows the structure of the tree T. The second axiom can be made more descriptive by characterising a tree-decomposition $(T, \mathcal{V})$ in terms of the separations of G induced by the edges of T: A pair $(T, \mathcal{V})$ is a tree-decomposition if and only if it satisfies (T1) and

(T2') For every edge $f = t_1t_2$ of T and for $i = 1, 2$, set $A_i^f \coloneqq \bigcup_{t \in V(T_i)} V_t$ where T_i is the component of $T - f$ which contains t_i. Then $s^f \coloneqq \{A_1^f, A_2^f\}$ 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 tree-decomposition $(T, \mathcal{V})$ of G is *into cliques* if all its bags are cliques in G. For the interested reader, we remark that a map which assigns to each vertex v of G a subgraph T_v of a tree T is a subtree representation of G if and only if $(T, \mathcal{V})$ is a tree-decomposition into cliques with $V_t \coloneqq \{v \in V(G) \mid t \in V(T_v)\}$. See [AK25, §2] for a more general correspondence.

A simple observation is that any finite clique is contained in a single bag of a tree-decomposition:

Lemma 3.2.6. *Let $(T, \mathcal{V})$ be a tree-decomposition of a (possibly infinite) graph G. For every finite[7] clique K in G, there is a node t of T such that the bag V_t contains K.*

Proof. Consider the subtrees T_v of T from (T2) with $v \in K$. By (T1), the T_v are pairwise intersecting. A folklore result asserts that every family of subtrees of a (possibly infinite) tree T is (finitely) Helly, i.e. whenever a finite family $\mathbf{T}$ of subtrees of T is pairwise intersecting, then their intersection $\bigcap \mathbf{T}$ is non-empty. Since K is finite, this yields that there is a node t that is contained in every T_v with $v \in K$, i.e. V_t contains K, as desired. □

3.2.3 Interplay of tree-decompositions and separations

First, 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 both $A \setminus B$ and $B \setminus A$ contain the vertex set of a full component of $G - (A \cap B)$. Analogous to separators, it is easy to see 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 $k \coloneqq |A \cap B|$ of its separator; we also say that $\{A, B\}$ is a *k-separation* of G.

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.

[7] For infinite cliques Lemma 3.2.6 fails trivially: A countably infinite clique $\{v_1, v_2, \dots\}$ has a tree-decomposition $\{T, \mathcal{V}\}$ into finite bags where $T = t_1t_2 \dots$ is a ray and $V_{t_i} = \{v_1, \dots, v_i\}$.

Recall that a tree-decomposition $(T, \mathcal{V})$ of a graph G is Γ-*canonical* for a subgroup Γ of $\mathrm{Aut}(G)$ if there is an action of Γ on T that commutes with $t \mapsto V_t$, i.e. $V_{\varphi \cdot t} = \varphi(V_t)$ for every $t \in V(T)$ and $\varphi \in \Gamma$; we will usually write $\varphi(t)$ instead of $\varphi \cdot t$ for notational simplicity. An $\mathrm{Aut}(G)$-canonical tree-decomposition is also called *canonical* for short. Furthermore, 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 3.2.7 ([DJKK26], 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 $t \mapsto V_t$.*

Given a tree-decomposition $\mathcal{T} := (T, \mathcal{V})$ of a graph G, it follows from the definition that the set $N(\mathcal{T}) := \{s^f : f \in E(T)\}$ of all its induced separations s^f of G is nested. If $\mathcal{T}$ is also Γ-canonical for a subgroup Γ of $\mathrm{Aut}(G)$, then the set $N(\mathcal{T})$ inherits this symmetry in that it is itself Γ-*invariant*: The automorphisms φ of G naturally act on the separations $s = \{A, B\}$ of G by the map $\varphi(s) := \{\varphi(A), \varphi(B)\}$. A set S of separations of G is Γ-*invariant*[8] if $\varphi(s) \in S$ for every $s \in S$ and every $\varphi \in \Gamma$.

Conversely, in finite graphs, a (Γ-invariant) nested set of separations of G induces a (Γ-canonical) tree-decomposition of G. In infinite graphs, this holds under an additional assumption, as shown in [EKT22b, Lemma 2.7]; we here state their lemma in a way compatible with [JKK25]. For this, let us call a set S of separations of G *point-finite* if every vertex of G appears in the separator of at most finitely many separations in S. Further, a tree-decomposition $(T, \mathcal{V})$ of G is *point-finite* if T_v is finite for every vertex v of G. Given a nested set N of proper separations of a graph, let $\mathcal{T}(N)$ be defined as in [JKK25, Construction 2.3].

Lemma 3.2.8 ([EKT22b, Lemma 2.7])**.** *Let N be a nested set of proper separations of a connected graph G. If N is point-finite, then $\mathcal{T}(N)$ is a tree-decomposition 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 Γ-invariant for a subgroup Γ of $\mathrm{Aut}(G)$, then $\mathcal{T}(N)$ is Γ-canonical.*

3.3 Canonical tree-decompositions into cliques

In this section, we prove Theorems 8 and 9.

3.3.1 Proof of Theorem 8

Let us restate Theorem 8:

Theorem 8. *Let G be a locally finite, connected graph. Then G is chordal if and only if G admits a canonical tree-decomposition into cliques.*

[8] Γ-invariant sets of separations are sometimes also called Γ-*canonical*, and $\mathrm{Aut}(G)$-invariant sets of separations are referred to as *canonical* for short (see e.g. [DJKK26]).

The backwards-implication is straightforward (Proposition 3.1.1). For the forwards-implication, we aim to find a nested set $N(G)$ of separations of G such that $\mathcal{T}(N(G))$ is the desired tree-decomposition into cliques.

For the construction of $N(G)$, we draw on the most recent developments in tangle-theory. We shall not need any formal details here, but use these results as a black-box. Let us give a rough sketch of our construction in the context of tangle-theory. Tangles are an abstract notion of highly connected substructures of graphs originating in the work of Robertson and Seymour on their Graph Minor Theorem [RS83–04]. One of the central results of tangle-theory is the so-called tangle-tree theorem which asserts that the relative position of the tangles of a graph to each other is tree-like and can be captured by a nested set of separations [RS83–04].

A *tree of tangles* is a nested set of separations with the property that every two tangles of the graph are efficiently distinguished by one of its elements; many modern constructions even yield *canonical* trees of tangles, that is, $\mathrm{Aut}(G)$-invariant such nested sets. The most recent methods in tangle-theory find trees of tangles by focusing on the sets of separations that efficiently distinguish two given tangles [CK24; EKT21; EKT22b; JKK25]. Such sets of efficient tangle-distinguishers can themselves be axiomatized in a similar manner as tangles, leading to the notion of *bottlenecks* [JKK25]. The quest to find a (canonical) tree of tangles of a graph then translates into finding a (canonical) nested set of separations which meets all of the graph's bottlenecks.

To find our desired nested set $N(G)$ of separations, we follow this approach: We will first show that the set of separations that efficiently distinguishes two maximal cliques in a locally finite, chordal graph is indeed a bottleneck. Secondly, we deduce that there is a canonical nested set $N(G)$ of separations which meets all such bottlenecks; this step draws on recent results about tree-of-tangles constructions via bottlenecks [JKK25]. Thirdly, we show that $N(G)$ induces a tree-decomposition of G and that this tree-decomposition is into cliques. Intuitively, if this tree-decomposition were not into cliques, then some part would contain two distinct maximal cliques and they must be efficiently distinguished by a separation in $N(G)$, which yields a contradiction.

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, $X \subseteq A$ and $Y \subseteq B$.[9] If $\{A, B\}$ has minimum order among the separations of G that separate X and Y, then $\{A, B\}$ distinguishes X and Y *efficiently*. Note that if $\{A, B\}$ distinguishes X and Y (efficiently), then so does its separator $A \cap B$. 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 3.3.1. *Let X and Y be two distinct maximal cliques in a locally finite, chordal graph G. Then X and Y are distinguished by a tight separation of order less than the sizes of X and Y.*

Proof. Since G is locally finite, the cliques X and Y are finite. Hence, Theorem 3.2.3 yields that an efficient X–Y distinguishing separation has order less than the size of X and Y.

[9] In this chapter, we always consider separations $\{A, B\}$ that efficiently distinguish two distinct maximal cliques X and Y, and by Lemma 3.3.1, such separations even satisfy that, up to renaming sides, $X \subseteq A$ and X meets $A \smallsetminus B$, and $Y \subseteq B$ and Y meets $B \smallsetminus A$.

By Corollary 3.2.5 and its proof, both the separator of the separation and the separation itself are tight. □

Given two maximal cliques X and Y in a locally finite, chordal graph G, let $\beta(X,Y) = \beta(Y,X)$ be the set of separations efficiently distinguishing X and Y; in particular, all separations in $\beta(X,Y)$ have the same finite order and, by Lemma 3.3.1, the separations in $\beta(X,Y)$ have order less than $|X|$ and $|Y|$. This implies by [JKK25, Example 3.4] that $\beta(X,Y)$ is a bottleneck. In this chapter, we will use the notion of bottlenecks as a black-box and refer the reader to [JKK25, Section 3] for the definition.

Lemma 3.3.2. *Let X and Y be two distinct maximal cliques in a locally finite, chordal graph G. Then the set $\beta(X,Y)$ is a bottleneck in G.*

Proof. Combine Lemma 3.3.1 with [JKK25, Example 3.4]. □

In a second step, we find a canonical nested set of separations which meets all the bottlenecks $\beta(X,Y)$. For this, we will apply [JKK25, Construction 3.6] to the set $\mathcal{B}$ of all such bottlenecks $\beta(X,Y)$. We restate this construction here:

Construction 3.3.3 ([JKK25, Construction 3.6])**.** Let G be a graph, and let $\mathcal{B}$ be a set of bottlenecks in G. For every $k \in \mathbb{N}$, let $\mathcal{B}^k(G)$ be the set of all k-bottlenecks in $\mathcal{B}$ in G and $S^k := \bigcup \mathcal{B}^k(G)$. Recursively for all $k \in \mathbb{N}$, define nested sets $N^k(G) \subseteq S^k$ of separations, as follows. Let $x^k \colon S^k \to \mathbb{N}$ assign to each separation $s \in S^k$ the number of separations in S^k that s crosses, assuming for now that this number is finite. 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 $N(G) := \bigcup_{k\in\mathbb{N}} N^k(G)$.

Construction 3.3.3 yields a well-defined nested set of separations, as shown in [JKK25]; in particular, the function $x^k : S^k \to \mathbb{N}$ is indeed well-defined.

Theorem 3.3.4 ([JKK25, Theorem 3.7])**.** *Let G be a locally finite, connected graph. Given any set $\mathcal{B}$ of bottlenecks in G, the sets $N(G)$ and $N^{<k}(G)$ (for $k \in \mathbb{N}$) as in Construction 3.3.3 are well-defined nested sets of tight separations, and $N^k(G,\beta)$ is non-empty for every $k \in \mathbb{N}$ and every k-bottleneck β in G. Moreover, if $\mathcal{B}$ is* $\mathrm{Aut}(G)$*-invariant, then so are $N(G)$ and $N^{<k}(G)$ for all $k \in \mathbb{N}$.*

Here, a set $\mathcal{B}$ of bottlenecks is $\mathrm{Aut}(G)$*-invariant* if the action of $\mathrm{Aut}(G)$ on the bottlenecks of G (which is inherited from the action of $\mathrm{Aut}(G)$ on the separations of G) is invariant on $\mathcal{B}$.

Corollary 3.3.5. *Let G be a locally finite, connected, chordal graph, and let $\mathcal{B}$ be the set of all bottlenecks in G of the form $\beta(X,Y)$ for any two distinct maximal cliques X and Y in G. Then $N(G)$ is a well-defined,* $\mathrm{Aut}(G)$*-invariant, nested set of tight separations of G that efficiently distinguishes every two distinct maximal cliques in G.*

Proof. By Lemma 3.3.2, the sets $\beta(X,Y)$ for distinct maximal cliques X and Y in G are bottlenecks, and their collection $\mathcal{B}$ is $\mathrm{Aut}(G)$-invariant by definition. By Theorem 3.3.4, the set $N(G)$ of separations obtained from Construction 3.3.3 is well-defined, $\mathrm{Aut}(G)$-invariant

and nested, and it meets every $\beta(X, Y)$. In particular, $N(G)$ efficiently distinguishes every two distinct maximal cliques in G. Moreover, each element of $N(G)$ is contained in some $\beta(X, Y)$ and hence tight by Lemma 3.3.2. □

The third and final step of our proof of Theorem 8 is to show that $N(G)$ induces a canonical tree-decomposition, for which we will apply Lemma 3.2.8 to obtain $\mathcal{T}(N(G))$, and that the canonical tree-decomposition $\mathcal{T}(N(G))$ is indeed into cliques. For later use, we will more generally show that these conclusions hold for any nested set N of tight separations of G which efficiently distinguishes every two distinct maximal cliques in G; by Theorem 3.3.4 this assumption in particular includes any set N coming from Construction 3.3.3 applied to a set $\mathcal{B}$ of bottlenecks containing all $\beta(X, Y)$.

We begin by checking that such a set N satisfies the assumption of Lemma 3.2.8, that is, N is point-finite.

Lemma 3.3.6. *Let G be a locally finite, connected chordal graph, and let N be a nested set of tight separations of G which efficiently distinguishes every two distinct maximal cliques in G. Then the 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, that is, N is point-finite. □

Theorem 3.3.7. *Let G be a connected, locally finite, chordal graph, and let N be an* ($\mathrm{Aut}(G)$*-invariant) nested set of tight separations of G that efficiently distinguishes every two distinct maximal cliques in G. Then $\mathcal{T}(N)$ is a (canonical) regular, point-finite tree-decomposition of G into cliques.*

Proof. By Lemma 3.3.6, N is point-finite, and the separations in N are proper since they are tight. So Lemma 3.2.8 yields 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. Note that, by Lemma 3.2.8, $\mathcal{T}$ is regular, point-finite and canonical.

It remains to show that $\mathcal{T}$ is into cliques. All adhesion sets of $\mathcal{T}$ are cliques since $N = \{s^f \mid f \in E(T)\}$ and the separators of separations in N are cliques by Lemma 3.3.6. Thus, every part $G_t := G[V_t]$ of $\mathcal{T}$ is connected, since G is connected. Moreover, as G is chordal, so are its induced subgraphs G_t.

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 $\mathcal{T}$. Now let X be an inclusion-minimal u–v separator in G_t. Then X is non-empty, as the part G_t of $\mathcal{T}$ is connected, and X is a clique by Theorem 3.2.1, as the induced subgraph G_t of G is chordal.

By applying Lemma 3.2.2 to the (full) components of $G_t - X$ containing u and v, respectively, 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. So since X is a clique, there are maximal cliques $X_{u'}$ and $X_{v'}$ of G that contain $X \cup \{u'\}$ and $X \cup \{v'\}$, respectively. It is easy to see that,

since 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 $X_{u'} \ni u'$ and $X_{v'} \ni v'$ in G. Every $X_{u'}$–$X_{v'}$ separator in G must contain their intersection $X_{u'} \cap X_{v'} \supseteq X$. Hence, the only $X_{u'}$–$X_{v'}$ separator of minimum size in G is X itself. Since $N = \{s^f \mid f \in E(T)\}$ efficiently distinguishes every two distinct maximal cliques such as $X_{u'}$ and $X_{v'}$ by assumption, some separation s^f corresponding to an edge f of T separates $X_{u'}$ and $X_{v'}$ efficiently. By the above argument, the separator of s^f is X; in particular, s^f separates $u' \in X_{u'} \setminus X$ and $v' \in X_{v'} \setminus X$. 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, which is a contradiction. □

Proof of Theorem 8. Apply Theorem 3.3.7 to the set $N(G)$ obtained from Construction 3.3.3 for the set $\mathcal{B}$ of all bottlenecks in G of the form $\beta(X, Y)$ for two distinct maximal cliques X and Y in G. By Corollary 3.3.5, this $N(G)$ satisfies the assumptions of Theorem 3.3.7. □

For experts, we remark that one may alternatively obtain a nested set like $N(G)$ by [EKT22b, Theorem 1.3], as every maximal clique in a chordal graph yields a 'robust, regular profile' and by invoking Lemma 3.3.1 one may prove that each two of them are 'distinguishable'. In contrast, the nested set from [DJKK26, Definition 5.2] does not work in general, as the 'tangles' induced by cliques X have in general only order $2/3 \cdot |X|$, which as a consequence may make the two tangles induced by two distinct maximal cliques indistinguishable. Finally, we can also not work with entanglements [CK24] as the order of the separations in clique-distinguishing bottlenecks is not finitely bounded, which would be a necessary assumption for their [CK24, Theorem 4.2].

3.3.2 Proof of Theorem 9

Theorem 8 asserts that every locally finite, chordal graph admits a canonical tree-decomposition into cliques. The bags of this decomposition can be non-maximal cliques, and there are even finite chordal graphs which do not have a canonical tree-decomposition into maximal cliques (see Example 3.5.1). In this section, we prove Theorem 9 which asserts that if the chordal graph $\hat{G}$ we seek to decompose is a normal covering of a locally finite graph G, then we can find a tree-decomposition into maximal cliques, which is still $\Gamma(p)$-canonical for the group $\Gamma(p)$ of the deck transformations of the normal covering $p : \hat{G} \to G$.

For this, recall that a *covering* of a (loopless) graph G is a surjective homomorphism $p\colon \hat{G} \to G$ such that p restricts to an isomorphism from the edges incident to any given vertex $\hat{v}$ of $\hat{G}$ to the edges incident to its projection $v := p_r(\hat{v})$. 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 $\Gamma(p)$. Further, a covering $p\colon \hat{G} \to G$ is *normal* if the group $\Gamma(p)$ of deck transformations of p acts transitively on the fibers $p^{-1}(v)$ of p. We remark that r-local coverings are normal, so the following theorem is applicable to r-local coverings and we will detail this application in Section 3.4.

Let us now restate Theorem 9, the main result of this section:

Theorem 9. *Let G be a locally finite, connected 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.*

Our plan to prove Theorem 9 is simple: We will take the canonical tree-decomposition $(T, \mathcal{V})$ of $\hat{G}$ into cliques as given by Theorem 8, and we then iteratively contract the $\Gamma(p)$-orbits of edges st of the decomposition tree T which witness that $(T, \mathcal{V})$ is not into maximal cliques, that is, with $V_s \subseteq V_t$. The difficult part of the proof is to show that this construction indeed results in a $\Gamma(p)$-canonical tree-decomposition into its maximal cliques.

One ingredient to our proof of Theorem 9 is a simple observation about the interaction of cliques in $\hat{G}$ with the group $\Gamma(p)$ of deck transformations of the covering p:

Lemma 3.3.8. *Let G be a connected graph and $p\colon \hat{G} \to G$ a covering. Then for every clique K in $\hat{G}$, 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 neither loops nor parallel edges. Thus, every two distinct elements of the same fiber of p have distance > 2 in $\hat{G}$. In particular, if two cliques in $\hat{G}$ meet, then their union cannot contain two distinct vertices of the same fiber of p. This immediately yields the statement, since $\Gamma(p)$ acts freely on $\hat{G}$. □

As a sidenote, we can deduce from Theorem 8 and Lemma 3.3.8 that the group of deck transformations of a normal covering is free if the cover is chordal:

Corollary 3.3.9. *Let G be a connected, locally finite graph and $p\colon \hat{G} \to G$ a covering. If $\hat{G}$ is chordal, then its group $\Gamma(p)$ of deck transformations is free.*

Proof. Set $\Gamma := \Gamma(p)$. By Theorem 8, $\hat{G}$ admits a Γ-canonical tree-decomposition $(T, \mathcal{V})$ into cliques. Lemma 3.3.8 ensures that Γ acts freely on the decomposition tree T. Hence, it follows from [L17b, Theorem 4.2.1] that Γ is free. □

Proof of Theorem 9. Let $(T, \mathcal{V})$ be any tree-decomposition of $\hat{G}$ that is point-finite, regular, $\Gamma(p)$-canonical and into cliques, such as the one given by Theorem 8.

Claim 1. *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 X of $\hat{G}$, there is a unique node $t(X)$ of T with $X = V_{t(X)}$, and*

(ii) *every V_t is a maximal clique.*

Proof of the claim. (i): First, we show that $t(X)$ exists. By Lemma 3.2.6, there exists a node t of T whose bag V_t contains X. As $(T, \mathcal{V})$ is into cliques and X is a maximal clique, V_t is indeed equal to X.

Secondly, we show that $t(X)$ is unique. Suppose for a contradiction that there are two distinct nodes t, t' of T with $V_t, V_{t'} = X$. Then (T2) ensures that all vertices s on the t–t' path in T satisfy $X \subseteq V_s$. As $(T, \mathcal{V})$ is into cliques and X is a maximal clique, we thus have $X = V_s$. Hence, every edge $f = ss'$ on the t–t' path in T contradicts the assumption, as $V_s = X = V_{s'}$.

(ii): Suppose for a contradiction that V_t is not a maximal clique. Let X be a maximal clique that contains V_t. Let $f = ts$ be the first edge on the t–$t(X)$ path in T. Then (T2) ensures that $V_t \subseteq V_s$, which contradicts our assumption. ■

We will now iteratively contract the $\Gamma(p)$-orbits of edges st of T with $V_s \subseteq V_t$, with the aim of obtaining a tree-decomposition satisfying the assumption of Claim 1. Note that since $(T, \mathcal{V})$ is $\Gamma(p)$-canonical and regular, the deck transformations of p actuniquely on T such that this action commutes with $t \mapsto V_t$, by Lemma 3.2.7.

Let us first note that this is a countable procedure in that T is countable and so is the number of $\Gamma(p)$-orbits of $E(T)$. Indeed, since $(T, \mathcal{V})$ is point-finite, all T_v are finite, and since the connected graph G is locally finite, G is countable. Hence, $T = \bigcup_{v \in V(G)} T_v$ is the union of countably many finite graphs, and thus countable as well.

Now pick an enumeration $M_1, M_2, \dots$ of the $\Gamma(p)$-orbits of $E(T)$. Set $I_0 := \emptyset$, and write $[i] := \{1, \dots, i\}$ for $i \in \mathbb{N}$. We iteratively construct an inclusion-wise increasing sequence of $I_i \subseteq [i]$ and consider tree-decompositions $(T^i, \mathcal{V}^i)$ of $\hat{G}$ obtained by contracting $M^i := \bigcup_{j \in I_i} M_j$ in T, i.e. set $T^i := T/M^i$, set $V^i_t := \bigcup_{s \in C} V_s$ if t is the contraction vertex given by some component C of $(V(T), M^i)$, and else $V^i_t := V_t$. The I_i will be constructed such that $(T^i, \mathcal{V}^i)$ 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 and $\Gamma(p)$-canonical, and that the $\Gamma(p)$-orbits of $E(T^i)$ are the M_j with $j \in [i] \setminus I_i$ and the M_j with $j > i$.

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 $I_i := I_{i-1}$ is as desired. Otherwise, we claim that $I_i := I_{i-1} \cup \{i\}$ is as desired. For this, we make the following observation, which is also the only part of the proof relying on the assumption that $\hat{G}$ is a covering of a graph G.

Claim 2. *The $\Gamma(p)$-orbit M of an edge st of T with $V_s \subseteq V_t$ is a matching in T.*

Proof of the claim. Suppose for a contradiction that M contains two incident edges m_1m_2 and m_2m_3 of T. Then there is a deck transformation $\gamma \in \Gamma(p) \setminus \{\mathrm{id}_{\hat{G}}\}$ that maps m_1m_2 to m_2m_3 (as unoriented edges). We cannot have $\gamma(m_2) = m_2$, since γ would then fix the clique V_{m_2} contradicting Lemma 3.3.8. So we must have $\gamma(m_1) = m_2$. But by our assumption on st, this implies either $V_{m_1} \subseteq V_{m_2} = \gamma(V_{m_1})$ or $V_{m_1} \supseteq V_{m_2} = \gamma(V_{m_1})$. In both cases, the image $\gamma(V_{m_1})$ of the clique V_{m_1} meets V_{m_1}, again contradicting Lemma 3.3.8. ■

Now 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}. By Claim 2, only independent edges st with, say, $V^{i-1}_s \subseteq V^{i-1}_t$ have been contracted, so $(T^i, \mathcal{V}^i)$ is still into cliques.

It remains to consider the limit step: Set $I := \bigcup_{i \in \mathbb{N}} I_i$, 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, $\Gamma(p)$-canonical, and we have $V'_s \not\subseteq V'_t$ and $V'_s \not\supseteq V'_t$ for every edge $st \in E(T')$. Finally, 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, contradicting the recursive construction. Altogether, $(T', \mathcal{V}')$ satisfies the assumptions of Claim 1, completing the proof. □

3.4 Canonical graph-decompositions via coverings for locally chordal graphs

In this section, we use our results on chordal graphs for the study of locally chordal graphs. With our applications of Theorems 8 and 9, we contribute to the extension of characterisations of chordal graphs to the locally chordal setting. Our results in particular give the first known characterisation of a graph class by a property of the canonical graph-decomposition $\mathcal{H}_r(G)$ that displays its r-global structure, as introduced in [DJKK26]:

Theorem 10. *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 $\mathcal{H}_r(G)$ of G which displays its r-global structure is into cliques.*

We give all definitions needed to understand and prove Theorem 10 in the course of this section.

3.4.1 Background: r-local coverings and r-locally chordal graphs

Let $r \in \mathbb{N}$. A (graph) 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 a (simple) graph G if and only if p is 1-ball preserving.

In [DJKK26] the *r-local covering* $p_r\colon G_r \to G$ of a connected graph G is introduced as the covering of G whose characteristic subgroup is the 'r-local subgroup' of the fundamental group of G. As the formal definition is not relevant to this chapter, we refer the reader to [DJKK26, §4] for details. In this chapter, we will often use that the r-local covering preserves the $r/2$-balls. Indeed, the r-local covering is the universal $r/2$-ball-preserving covering of G [DJKK26, Lemma 4.2, 4.3 & 4.4], 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 $p_r = p \circ q$. Following this equivalent description of the r-local covering, it is also defined for non-connected graphs G.

We refer to the graph G_r as the *r-local cover* of G. Note that the 0-, 1- and 2-local covers of a (simple) graph G are forests. For a reader familiar with coverings, we remark that the 0-, 1- and 2-local coverings are all indeed the universal covering of G.

A graph is *r-locally chordal* if its balls of radius $r/2$ are chordal. The following theorem gives various characterisations of r-locally chordal graphs; for our applications, the equivalence between (i) and (iii) is the relevant one.

Theorem 3.4.1 ([AKK25, Theorem 1])**.** *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) *G is r-chordal and wheel-free.*

(iii) *The r-local cover G_r of G is chordal.*

(iv) *Every minimal r-local separator of G is a clique.*

3.4.2 Background: Graph-decompositions and coverings

Graph-decompositions, as introduced in [DJKK26], generalize the notion of tree-decompositions by allowing the bags V_h of a decomposition $(H, \mathcal{V})$ to be arranged along a general graph H instead of a tree.

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) for every vertex v of G, the subgraph $H[W_v]$ of H induced on the set W_v of nodes whose bags contain v is connected.

We refer to H as the *model graph*, and also say that $(H, \mathcal{V})$ is a *graph-decomposition over* or *modelled on* H. The vertex-subsets V_h are called the *bags* of $\mathcal{H}$, and the W_v are called the *co-bags* of $\mathcal{H}$. One may flesh out the bags and co-bags by choosing spanning subgraphs G_h of $G[V_h]$ with $G = \bigcup_{h \in H} G_h$ and connected, spanning subgraphs H_v of $H[W_v]$. Then the G_h are the *parts* and the H_v are the *co-parts* of the decomposition $\mathcal{H}$.

Given a normal covering $p\colon \hat{G} \to G$ of a graph G and a $\Gamma(p)$-canonical tree-decomposition of $\hat{G}$, we can obtain a graph-decomposition of G by 'folding' the tree-decomposition along p, as described by the action of $\Gamma(p)$:

Construction 3.4.2 ([DJKK26, Construction 3.8 and Lemma 3.9])**.** Let $p\colon \hat{G} \to G$ be any normal covering of a (possibly infinite) connected graph G. Let $(T, \hat{\mathcal{V}})$ be a Γ-canonical tree-decomposition of $\hat{G}$, where $\Gamma := \Gamma(p)$ is the group of deck transformations of p. We define the graph-decomposition $(H, \mathcal{V})$ *obtained from* $(T, \hat{\mathcal{V}})$ *by folding via* p as follows. The model graph H is the orbit-graph of T under the action of Γ, i.e. $H = T/\Gamma$. Denote the quotient map from T to the orbit-graph $H = T/\Gamma$ also by p. The bag V_h corresponding to a node h of H is $p(\hat{V}_t)$ for a vertex (equivalently: every vertex) t in the Γ-orbit h. The parts G_h of $(H, \mathcal{V})$ are the projections $p(\hat{G}[\hat{V}_t])$. As co-parts H_v of $(H, \mathcal{V})$ we choose the projections $p(T_{\hat{v}})$ of the co-part $T_{\hat{v}}$ of $\mathcal{T}$ corresponding to a (equivalently: every) lift $\hat{v}$ of v to $\hat{G}$, where we denote also the quotient map from T to the orbit-graph $H = T/\Gamma$ by p.

3.4.3 Coverings and Cliques

In this section, we study the interactions of an r-local covering $p_r : G_r \to G$ with the (maximal) cliques of G_r and G. Some results in this section also appear in [AK25] taking a slightly different perspective (see [AK25, Lemma 4.8] and its proof, and our proof of Lemma 3.4.3 (ii)).

Lemma 3.4.3. *Let G be a (possibly infinite) graph and $r \geqslant 3$ an integer. Let $\mathcal{T}$ be a $\Gamma(p_r)$-canonical tree-decomposition of G_r, and let $\mathcal{H}$ be the graph-decomposition of G obtained from $\mathcal{T}$ by folding via p_r.*

(i) *$\mathcal{T}$ is into cliques if and only if $\mathcal{H}$ is into cliques.*

(ii) *$\mathcal{T}$ is into maximal cliques if and only if $\mathcal{H}$ is into maximal cliques.*

As tools for the proof of Lemma 3.4.3, we have some easy observations about cliques in covers and a lemma about tree-decompositions into disjoint unions of cliques.

Lemma 3.4.4 ([JKK25, Lemma 5.15])**.** *Let G be a (possibly infinite) graph and let $r \geqslant 3$ an integer.*

(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 $X \subseteq V(G)$ and every lift $\hat{x}$ of some vertex $x \in X$ to G_r, there exists a unique clique $\hat{X}$ of G_r containing $\hat{x}$ such that p_r restricts to a bijection from $\hat{X}$ to X.*[10]

(iii) *$N_{G_r}[\hat{X}] \cap N_{G_r}[\gamma(\hat{X})] = \emptyset$ for every clique $\hat{X}$ of G_r and every $\gamma \in \Gamma(p_r) \setminus \{\mathrm{id}_{G_r}\}$.*

Lemma 3.4.5. *If $(T, \mathcal{V})$ is a tree-decomposition of a connected graph G such that each $G[V_t]$ is a disjoint union of complete graphs, then $(T, \mathcal{V})$ is into cliques, i.e. each $G[V_t]$ is in fact a complete graph.*

Proof. Suppose for a contradiction that there is $t \in V(T)$ such that $G[V_t]$ is the disjoint union of at least two complete graphs; in particular, $G[V_t]$ has at least two connected components C and C'. Since G is connected, there is a path connecting such $V(C)$ and $V(C')$ in G, and we choose P to be a shortest such path. Then P meets every clique in G in at most a vertex or an edge.

Let T' be a node-minimal subtree of T containing t such that $V(P) \subseteq \bigcup_{s \in T'} V_s$. Since P meets two distinct components of $G[V_t]$, it follows that $T' \setminus \{t\}$ is non-empty. Let $t' \neq t$ be any leaf of T'. Then $P \cap G[V_{t'}]$ is a disjoint union of subpaths P' of P, each of length at most 1. Now an end vertex of P' is either also an end vertex of P or it has a neighbour outside of $V_{t'}$. In either case, the end vertex of P' must be contained in $V_{t'} \cap V_{t''}$, where t'' is the unique neighbour of the leaf t' of T. Thus, the path P', having length at most 1, is also contained in $G[V_{t''}]$. This implies $V(P) \subseteq \bigcup_{s \in T' \setminus \{t'\}} V_s$, contradicting the choice of T'. □

Proof of Lemma 3.4.3. (i): If the tree-decomposition $\mathcal{T}$ is into cliques, then Lemma 3.4.4 (i) yields that the graph-decomposition $\mathcal{H}$ is into cliques. Conversely, suppose that the graph-decomposition $\mathcal{H}$ is into cliques. We claim that the tree-decomposition $\mathcal{T} =: (T, \mathcal{V})$ of G_r has the property that each of its bags is a disjoint union of cliques; Lemma 3.4.5 then completes the proof. Since $\mathcal{H}$ is into cliques, p_r maps each bag V_t to a clique X of G. By Lemma 3.4.4 (ii) and (iii), the cliques $\hat{X}$ of G_r for which p_r restricts to a bijection to X are precisely the vertex sets of the connected components of $G_r[p_r^{-1}(X)]$. As $V_t \subseteq p_r^{-1}(X)$, every connected component of $G_r[V_t]$ is hence complete, as desired.

[10] This statement is slightly stronger than what is stated in [JKK25, Lemma 5.15], but follows immediately from its proof.

(ii): Given a maximal clique $\hat{X}$ of G_r, Lemma 3.4.4 (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 3.4.4 (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 may assume that $\hat{Y}$ meets $\hat{X}$. Since distinct lifts of the clique X are disjoint by Lemma 3.4.4 (iii), $\hat{Y}$ must contain the whole lift $\hat{X}$ of X. In particular, since $\hat{X}$ is a maximal clique of $\hat{G}$, $\hat{Y} = \hat{X}$, and thus $X = Y$ is a maximal clique of G. Similarly, one proves that every clique of G_r that projects to a maximal clique of G is also maximal. So since distinct cliques of G_r that project to the same clique of G are disjoint by Lemma 3.4.4 (iii), $\mathcal{T}$ is into maximal cliques if and only if $\mathcal{H}$ is into maximal cliques. □

We remark that both Lemmas 3.4.3 and 3.4.4 and their proofs hold more generally for all coverings $p : \hat{G} \to G$ which are 3/2-ball preserving.

3.4.4 Canonical graph-decompositions into cliques

Recall the theorem giving the canonical graph-decomposition $\mathcal{H}_r(G)$:

Theorem 3.1.3. ([DJKK26, Theorem 1]) *Let G be a finite graph, and $r > 0$ an integer. Then G has a unique canonical decomposition $\mathcal{H}_r(G)$ modelled on another finite graph $H = H(G, r)$ that displays its r-global structure.*

To understand Theorem 3.1.3, let us detail the construction of the graph-decomposition $\mathcal{H}_r(G)$ given by Theorem 3.1.3. We will not give the original construction from [DJKK26], but the most recent one from [JKK25], which builds on the notion of bottlenecks instead of tangles[11].

Let $p_r : G_r \to G$ be the r-local covering of a locally finite, connected, r-locally chordal graph G. By Theorem 3.4.1, G_r is a locally finite, connected and chordal graph. Let $M(G_r)$ be the set of separations obtained by applying Construction 3.3.3 to the set $\mathcal{B}$ of all bottlenecks in G_r. By Theorem 3.3.4, $M(G_r)$ is an $\mathrm{Aut}(G_r)$-invariant nested set of tight separations, and it in particular efficiently distinguishes every two maximal cliques in G. Thus, we can apply Theorem 3.3.7 to find that $\mathcal{T}_r \coloneqq \mathcal{T}(M(G_r))$ is canonical, regular, point-finite tree-decomposition of G_r. We can thus invoke Construction 3.4.2 to obtain a canonical graph-decomposition $\mathcal{H}(G, r)$ of G. For a thorough discussion on how $\mathcal{H}(G, r)$ displays the r-global structure of G and why it is the unique canonical such graph-decomposition of G, we refer the reader to [DJKK26].

Proof of Theorem 10. First, assume that G_r is chordal. By Theorem 3.3.7, $\mathcal{T}(M(G_r))$ is into cliques, and by Lemma 3.4.3 (i), $\mathcal{H}_r(G)$ is thus also into cliques.

Conversely, assume that $\mathcal{H}_r(G)$ is into cliques. Then Lemma 3.4.3 (i) yields that the tree-decomposition $\mathcal{T}(M(G_r))$ of G from which $\mathcal{H}_r(G)$ is obtained by folding via p_r is into cliques. Hence, this tree-decomposition witnesses by Proposition 3.1.1 that the r-local cover G_r of G is chordal. □

[11] As described in Section 3.3.1, sets of separations efficiently distinguishing two tangles form a bottleneck. Thus, the approach building on bottlenecks captures more than a tree of tangles. We refer the reader to [JKK25], in particular §1.4, for a detailed comparison of the two constructions.

3.4.5 Graph-Decompositions into maximal cliques

In this section, we discuss Theorem 9 in the context of r-locally chordal graphs. Its main application lies in the proof of a further characterization of r-locally chordal graphs; for this, a graph-decomposition$(H, \mathcal{V})$ of a graph G is *r-acyclic* for an integer $r \geqslant 0$ if for every set X of at most r vertices of G, the union of the respective co-parts H_x with $x \in X$ is acyclic.

Theorem 3.4.6. *[AK25, Theorem 4] 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 admits an r-acyclic graph-decomposition into maximal cliques.*

(iii) *G admits a canonical r-acyclic graph-decomposition into cliques.*

(iv) *G admits an r-acyclic graph-decomposition into cliques.*

For the proof of Theorem 3.4.6, we refer the reader to [AK25]. Theorem 8 features in the proof of '(i) implies (iii)' in that it gives the input to Construction 3.4.2, which then yields the desired canonical graph-decompositions into cliques by Lemma 3.4.3 (i) (as in the proof of Theorem 10). Similarly, Theorem 9 features in the proof of '(i) implies (ii)' in that it gives the input to Construction 3.4.2, which then yields the desired graph-decompositions into maximal cliques by Lemma 3.4.3 (ii). As shown in [AK25], both these graph-decompositions turn out to be also r-acyclic.

3.5 Counterexamples to possible strengthenings of Theorem 8

Theorem 8 asserts that every locally finite, chordal graph admits a canonical tree-decomposition into cliques. In this section, we provide counterexamples to two possible strengthenings of Theorem 8. First, Example 3.5.1 shows that we cannot obtain a canonical tree-decompositions into maximal cliques. Secondly, Example 3.5.2 demonstrates that a (countable) graph need not admit a canonical tree-decomposition into cliques, even if it admits non-canonical one.

3.5.1 Canonical tree-decompositions into maximal cliques

Recall that Halin's tree-decomposition in Theorem 3.1.2 are not only into cliques but even into maximal cliques. Formally, a tree-decomposition $(T, \mathcal{V})$ of a graph G is *into its maximal cliques* if the map $t \mapsto V_t$ is a bijection from $V(T)$ to the set of maximal cliques of G. It is natural to ask if Theorem 8 extends to canonical tree-decompositions into maximal cliques. But there are even finite (connected) chordal graphs that do not admit a canonical tree-decomposition into their maximal cliques:

Example 3.5.1. 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.

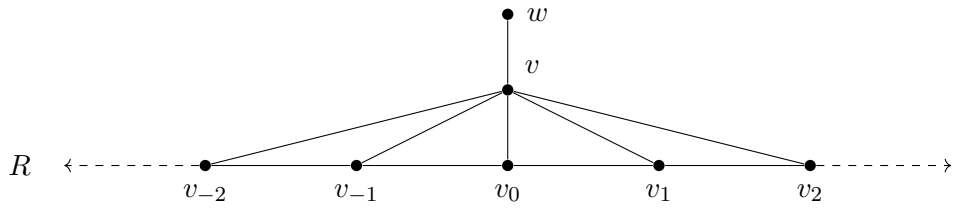

Figure 3.1: The graph G for Example 3.5.2 is obtained from a double-ray $R = \dots e_{-1}v_{-1}e_0v_0e_1v_1\dots$ by adding an apex vertex v and a further edge $e = vw$ attached to v.

Proof. For every tree-decomposition $(T, \mathcal{V})$ of $K_{1,t}$ into its maximal cliques, 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 decomposition graph 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, if $t \geqslant 3$, then the graph $K_{1,t}$ does not admit a canonical tree-decomposition into its maximal cliques. □

As a sanity check, we remark that the stars $K_{1,t}$ with $t \geqslant 3$ from Example 3.5.1, however, admit a canonical tree-decomposition into cliques: take a $K_{1,t}$ as the decomposition tree and associate with each of its leaves the incident edge as its bag and for the center vertex, choose the bag as the singleton center vertex, which is a non-maximal clique.

Further, we remark that Example 3.5.1 is still a counterexample if we drop the requirement of bijectivity from the definition of 'into its maximal cliques', that is, if we allow the tree-decomposition to have the same maximal clique as different bags. Indeed, the nodes of the decomposition tree T with the same bag form a subtree of T, and the edge-transitivity of $K_{1,t}$ then translates into transitivity of all those subtrees of T. So after contracting each such subtree to a single node, we can conclude with the same argument as above.

In light of Diestel's characterisation of countable graphs that admit a tree-decomposition into cliques via forbidden simplicial minors [D88], it might be a natural goal to characterise the countable, chordal graphs that admit canonical tree-decompositions into cliques or characterise the countable, chordal graphs that admit tree-decompositions into its maximal cliques.

3.5.2 Canonical tree-decomposition into cliques

Does every graph admitting a tree-decomposition into cliques also admit one that is canonical? We answer this question in the negative, even for countable graphs. Our example is inspired by [CHM22, Figure 1], who attribute their example to [DK15].

Example 3.5.2. The graph G depicted in Figure 3.1 is a countable, chordal graph that admits a tree-decomposition into cliques but not a canonical such tree-decomposition.

Proof. For a tree-decomposition into cliques, we take the double-ray R together with the edge v_0v as the decomposition tree. We let $V_v = \{v, w\}$ and $V_{v_i} = \{v, v_i, v_{i+1}\}$. It is easy

to check that this is indeed a tree-decomposition, and it is by the definition of G one into cliques.

Suppose now that $(T, \mathcal{V})$ is a tree-decomposition of G into cliques which is even canonical. The two triangles $v_{-1}v_0v$ and v_0v_1v are both contained in bags of this tree-decomposition, and these bags have to be distinct because $(T, \mathcal{V})$ is into cliques. Further, these two bags intersect in $\{v_0, v\}$, hence there is an edge $e_0 \in T$ with precisely $\{v_0, v\}$ as its adhesion set. Let φ be the (unique) automorphism of G which translates the ray R by one, i.e. $\varphi(e_i) = e_{i+1}$ and $\varphi(v_i) = \varphi(v_{i+1})$ for every $i \in \mathbb{Z}$, while $\varphi(v) = v$ and $\varphi(w) = w$. Repeated application of φ (and φ^{-1}) implies together with the canonicity of $(T, \mathcal{V})$ that e_0 yields the existence of edges e_i for all i with adhesion set $\{v_i, v\}$. The separation induced by e_i contains the edge $\{v, w\}$ on one side, and since $(T, \mathcal{V})$ is canonical, all separations induced by some e_i must agree on this side as the repeated application of φ (and φ^{-1}) shows. But since $(T, \mathcal{V})$ is a tree-decomposition, this is not possible. □

3.5.3 Tree-decompositions into maximal cliques

Does every graph admitting a tree-decomposition into cliques also admit one into its maximal cliques? We answer this question in the negative, even for countable graphs. We will deduce this from the fact that a certain uncountable chordal graph G^* does not admit a tree-decomposition into cliques, as Diestel showed in [D91, Paragraphs after Theorem 1.4]. A tree T at a vertex r naturally induces a partial order, the *tree-order* $\leqslant_T$, $u \leqslant_T v$ on its vertices by u lies on the unique r–v path rTv in T.

Example 3.5.3. Let T_2 be the infinite binary tree with root r, i.e. the root r has degree 2 and every other vertex of T_2 has degree 3. We obtain a chordal graph G from T_2 by adding for each two $\leqslant_{T_2}$-comparable vertices $u \leqslant v$ of T_2 the edge uv to G. Then G admits a (canonical) tree-decomposition into cliques, but it does not admit one into its maximal cliques.

Proof. For a tree-decomposition of G into cliques, we choose $T = T_2$ as the decomposition graph and for each $t \in T$, we let $V_t = \{v \in G : v \leqslant_{T_2} t\}$, the vertex set of the unique r–t path in T_2. It can be easily checked that this is indeed a tree-decomposition, and its bags are cliques by the definition of G from T_2.

Suppose for a contradiction that G admits a tree-decomposition, $(T, \mathcal{V})$ into maximal cliques. Let G^* be the graph obtained from G by adding to each ray R in the tree T_2 starting from r a vertex v_R and an edge uv_R to each vertex u of R. Note that the maximal cliques of G are precisely the vertex sets of the rays in T_2 starting from r. Thus, for each such ray R, we add v_R to the bag V_{t_R} to obtain a tree-decomposition $(T, \mathcal{V}')$ of G^* which is not only into cliques but even into maximal cliques. This contradicts that G^* does not admit a tree-decomposition into cliques, as Diestel showed in [D91, Paragraphs after Theorem 1.4]. □

PART II

Tangles and concrete clusters

CHAPTER 4

On vertex sets inducing tangles

This chapter is based on [ABJKW24] and joint work with Sandra Albrechtsen, Hanno von Bergen, Paul Knappe and Paul Wollan.

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.

4.1 Introduction

4.1.1 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 [RS83–04]. 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 tangles [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 4.1.1. *[DHL19] Is every tangle in a graph G induced by some set $X \subseteq V(G)$?*

In what follows we will often consider Problem 4.1.1 for all tangles of some fixed order $k \in \mathbb{N}$ and then say for short: Problem 4.1.1 for k.

We remark that Problem 4.1.1 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 [E15, Theorem 1.2] proved Problem 4.1.1 directly. Independently, Grohe [G16, 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 4.1.1 for $k = 3$. Similarly, it should be possible to derive Problem 4.1.1 for $k = 4$ from the recent characterisation of 4-tangles in terms of 'internal 4-connectedness' by Carmesin and Kurkofka [CK25, Theorem 1]. Besides these results for small k, Diestel, Elbracht and Jacobs [DEJ24, Theorem 12] showed that Problem 4.1.1 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 4.1.1 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 4.1.2. *[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 4.1.1, but also Theorem 4.1.2 fails for tangles in general discrete contexts, such as matroids or data sets (see e.g [D18a; DHL19; DO19; EKT21]). However, no such example is known for tangles in graphs. Thus, Problem 4.1.1 is open for all $k \geqslant 4$.

4.1.2 Our contributions to Problem 4.1.1

In this chapter, we reduce Problem 4.1.1 for every k to graphs whose size is bounded by a function in k:

Theorem 11. *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.*

[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.

As an immediate corollary of Theorem 11, we obtain that one may verify the validity of Problem 4.1.1 for any fixed k computationally in explicitly bounded, though impractically long, time by checking Problem 4.1.1 for every k-tangle in every connected graph with fewer than M edges.

Corollary 12. *For $k \geqslant 1$, there exists $M = M(k) \in O(3^{k^{k^5}})$ such that Problem 4.1.1 holds for k if it holds for all k-tangles in connected graphs G with fewer than M edges.*

Our second corollary of Theorem 11 asserts that whenever Problem 4.1.1 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 4.1.1 that induces a given k-tangle with $k \leqslant 2$ to be of size at most k; this follows from the characterisation of 1- and 2-tangles by components and blocks, respectively. More generally, we prove that for every k-tangle τ in a graph G one may choose a weight function w that induces τ, which exists by Theorem 4.1.2, in such way that its *total weight* $w(V(G)) = \sum_{v \in V(G)} w(v)$, which it distributes on $V(G)$, is bounded in k.

Corollary 13. *For every integer $k \geqslant 1$, there exists $K = K(k)$ such that for every k-tangle τ in a graph G there exists a weight function $w : V(G) \to \mathbb{N}$ which induces τ and whose total weight $w(V(G))$ is bounded by K. In particular, the support of w has size $\leqslant K$.*

Moreover, if Problem 4.1.1 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 11.

We derive Corollary 13 from Theorem 11 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 4.1.2) 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 11 by fixing each such weight function to be an indicator function of some inducing set given by the assumed positive answer to Problem 4.1.1 (see Section 4.8 for details).

4.1.3 An inductive proof method for tangles

Our proof of Theorem 11 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 14. *For every integer $k \geqslant 1$ there is some $M(k) \in O(3^{k^{k^5}})$ such that the following holds: Let τ be a k-tangle in a graph G. Then there exists a sequence $G_0, \ldots, G_m$ of graphs and k-tangles τ_i in G_i for every $i \in \{0, \ldots, m\}$ such that*

- *$G_0 = G$, $\tau_0 = \tau$;*
- *G_i is obtained from G_{i-1} by deleting an edge, suppressing a vertex, or taking a proper component;*
- *the k-tangle τ_{i-1} in G_{i-1} survives as the k-tangle τ_i in G_i for every $i \in [m]$;*
- *G_m is connected and has fewer than $M(k)$ edges.*

Before we proceed to explain the crucial term 'survive' in this theorem, we remark that it is necessary to allow the suppression of a vertex in Theorem 14. 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 4.3.6).

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 $A \setminus B$ and $B \setminus A$. A k-tangle τ in G *extends* to a k-tangle τ' in G' if $\tau \subseteq \tau'$. Then we also say that the k-tangle τ *survives* as the k-tangle τ' in G'. We remark that such an extension τ' of τ may or may not exist in G'; if it exists, it need not be unique. If such a τ' exists and is unique, then we also say that τ *induces* τ'.

For the cases that G' is a component of G or obtained from G by suppressing a vertex of G, we refer the reader to Section 4.3.2 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 [CK25, Lemma 2.1]), we remark that 'surviving' is the reverse notion.

4.1.4 Overview of the proof of Theorem 14

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 4.3.1), and

(ii) if $k = 2$, then τ extends to a k-tangle in $G - e$ for some edge e of G (Lemma 4.3.2).

A rather simple analysis will yield that

(iii) if G is not connected, then τ induces a k-tangle in a unique component of G (Lemma 4.3.3),

(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 4.3.4), 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 4.3.5).

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 4.3). 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 15. *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 15 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.

4.1.5 Proof sketch of Theorem 15

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 4.4). 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 4.4.3). 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 4.4.4).

Second, we assume that the graph G contains no tangle of order $> k$ (Section 4.7). 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 4.6). To obtain the desired decomposition, we start with the tree-decomposition obtained from the tangle-tree duality theorem [D17b, 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 4.5.1).

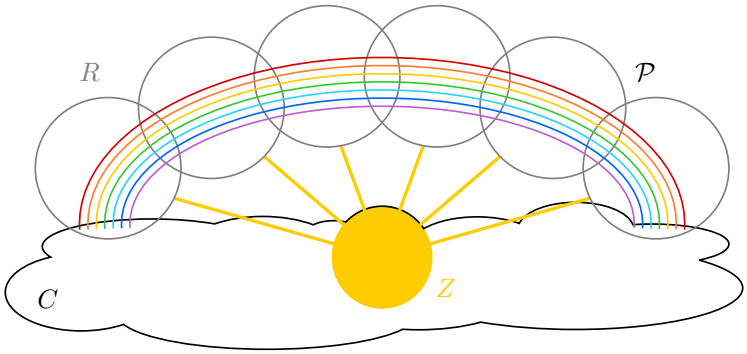

Figure 4.1: Schematic drawing of a rainbow-cloud decomposition

Roughly speaking (see also Figure 4.1), 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 4.5 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 4.7.1).

4.1.6 How this chapter is organised

We recall the relevant terminology concerning tangles in Section 4.2. We then introduce in Section 4.3 the notion of 'survives' and prove Theorem 14 in several special cases: the case $k \leqslant 2$, the case $\delta(G) \leqslant 2$, and when the graph G is disconnected. In Section 4.4 we then prove Theorem 15, the remaining case of Theorem 14 when G contains a tangle of order $k+1$. In Section 4.5, 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 4.6 about the interaction of rainbow-cloud decompositions and separations, we complete the proof of Theorem 15 in Section 4.7. In Section 4.8 we collect all the individual cases to prove Theorem 14, derive Theorem 11, and deduce Corollary 12 and Corollary 13.

4.2 Preliminaries

In this section, we recall the relevant terms and definitions concerning tangles in graphs. For general graph-theoretic concepts and notation, we follow [D17b]. 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.

4.2.1 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 $A \setminus B$ and $B \setminus A$ 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 [B48, 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 $< k$.

4.2.2 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:

Proposition 4.2.1. *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

[2]Recall that a *block* of a graph is a $\subseteq$-maximal connected subgraph H such that $H - v$ is connected for every vertex $v \in H$. In particular, as the empty graph is not considered to be connected, every block contains some edge.

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.

4.2.3 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.

4.2.4 Lifts of tangles

It is well-known that tangles in minors lift to tangles in the host graph (see [RS91, (6.1)] or [CK25, 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 $i = 1,2$.

(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.

4.3 Definition of 'survive' and special cases of Theorem 14

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 4.3.1 and then for graphs G' obtained from G by suppressing a single vertex of degree 2 in Section 4.3.2. Alongside these definitions, we prove several lemmas which all deal with special cases of Theorem 14. Finally, we provide Example 4.3.6 which demonstrates that suppressing vertices of degree 2 needs to be allowed in Theorem 14.

4.3.1 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 4.3.1. *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 4.2.1, 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 4.2.1 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 4.3.1 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 4.3.2. *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 4.2.1, $G[X_\tau]$ 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 4.2.1, 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 4.3.1 also in the proof of Lemma 4.3.2 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 4.3.3. *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 4.2.1, $G[X]$ 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 4.3.4. *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. □

4.3.2 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 4.3.5 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 4.3.5. *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. □

4.3.3 Suppressing vertices of degree 2 in Theorem 14

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 14 it is necessary to allow the suppression of a vertex.

Example 4.3.6. 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 4.3.3, 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 4.3.4 and 4.3.5 where $e' := e$ if $e \in E(G)$ or $e' := uv$ otherwise. □

4.4 Proof of Theorem 15 if G has a higher-order tangle

In this section we prove Theorem 15 for graphs which contain a tangle of order $> k$.

Theorem 4.4.1. *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 4.4.3). 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 4.4.4).

For both cases, we will make use of the following observation.

Lemma 4.4.2. *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 4.4.3. *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 4.4.2. 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 4.4.3 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 4.4.2. 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 4.4.4. *Let $k \geqslant 2$ be an integer, and let G be a graph with a k-tangle τ. If there exists a $(k+1)$-tangle $\tilde{\tau}$ in G with $\tau \not\subseteq \tilde{\tau}$, then there exists an edge $e \in E(G)$ such that τ extends to a k-tangle τ' in $G - e$.*

Proof. We first find an edge e of G that we afterwards prove to be as desired. For this, let $(B, A) \in \tilde{\tau}$ be a separation of G which distinguishes τ and $\tilde{\tau}$ and is $\leqslant$-maximal in $\tilde{\tau}$ among all such distinguishing separations. Note that (B, A) is even maximal in $\tilde{\tau} \cap \vec{S}_k$: any separation $(C, D) \in \tilde{\tau} \cap \vec{S}_k$ with $(B, A) < (C, D)$ would also distinguish τ and $\tilde{\tau}$ since $(D, C) \in \tau$ by the consistency of τ. We then choose an arbitrary edge e in $G[A \setminus B]$. Let us first show its existence.

There exists a vertex $v \in A \setminus B$, since otherwise $(V(G), A) = (B, A) \in \tilde{\tau}$, contradicting that $\tilde{\tau}$ is a tangle. If the vertex v has a neighbour in $A \setminus B$, then the edge joining them is as desired. So suppose for a contradiction that all neighbours of v are in $A \cap B$. We can then find a forbidden triple in $\tilde{\tau}$: First, we can move v from $A \setminus B$ to $B \setminus A$ to obtain a new separation $\{A \setminus \{v\}, B \cup \{v\}\}$ of G, which has the same order as $\{A, B\}$. Thus, $\tilde{\tau}$ contains an orientation of it, and we must have $(A \setminus \{v\}, B \cup \{v\}) \in \tilde{\tau}$ due to the maximality of (B, A) in $\tilde{\tau} \cap \vec{S}_k(G)$. Secondly, since $|A \cap B| < k$ and $\tilde{\tau}$ is a $(k+1)$-tangle in G, we have $((A \cap B) \cup \{v\}, V(G)) \in \tilde{\tau}$. Hence, $\{(B, A), (A \setminus \{v\}, B \cup \{v\}), ((A \cap B) \cup \{v\}, V(G))\}$ is contained in the tangle $\tilde{\tau}$, but it is also a forbidden triple, which is a contradiction. All in all, v has a neighbour in $A \setminus B$; in particular, $G[A \setminus B]$ contains an edge.

From now on, we prove that the chosen edge $e \in G[A \setminus B]$ is as desired. For this, we consider $G' := G - e$ and construct an orientation τ' of $S_k(G')$ to which τ extends. It then remains to show that τ' is a tangle in G'.

For the construction of τ', note that τ' has to contain not only τ, but also all orientations of separations of $S_k(G')$ that are 'forced' by the request that τ' shall again be a tangle and hence especially consistent. More formally, we say that τ *forces* an orientation of a separation $\{C, D\}$ of G' if there exists a separation $(E, F) \in \tau$ such that $(C, D) \leqslant (E, F)$ or $(D, C) \leqslant (E, F)$. In particular, τ forces an orientation of every separation in $S_k(G')$ which is also a separation of G and the separations in $\vec{S}_k(G')$ which are maximal among all those forced by τ are separations of G. Note that the consistency of τ ensures that at most one orientation of a separation in $S_k(G')$ is forced by τ.

We now define the orientation τ' of $S_k(G')$. If τ forces an orientation of a separation $\{C, D\} \in S_k(G')$, then we put the respective orientation in τ'. Otherwise, $\{C, D\}$ is especially not a separation of G, so e has one endvertex in $C \setminus D$ and the other one in $D \setminus C$. Then $\{C \cup e, D\}$ and $\{C, D \cup e\}$ are separations of G of order at most k, as $\{C, D\}$ has order less than k. Thus, both these separations are oriented by the $(k+1)$-tangle $\tilde{\tau}$ in G, and by Lemma 4.4.2, 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 of the claim. 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 of the claim. Since $e \in G[A \setminus B]$, $\{B \cap C, A \cup D\}$ is a separation of G. Assume that it has order less than k. Then τ contains an orientation of $\{B \cap C, A \cup D\}$, and this orientation must not be $(A \cup D, B \cap C)$ as τ would otherwise force the orientation (D, C) of $\{C, D\}$ to be in the orientation τ' which already contains (C, D).

It remains to show that $\{B \cap C, A \cup D\}$ indeed has order less than k; suppose for a contradiction otherwise. Since both $\{A, B\}$ and $\{C, D\}$ have order at most $k - 1$, this implies that $\{B \cup C, A \cap D\}$ has order less than $k - 1$ by submodularity. The edge e is in $G[A \setminus B]$ by its choice. Additionally, it has one endvertex in $C \setminus D$ and the other one in $D \setminus C$ because $\{C, D\}$ is not a separation of G by assumption. Therefore, the order of $\{B \cup C \cup e, A \cap D\}$ increases compared to the order of $\{B \cup C, A \cap D\}$ by exactly one. So $\{B \cup C \cup e, A \cap D\}$ has order $< k$. We now show that none of its orientations can be contained in the tangle $\tilde{\tau}$, which then yields the desired contradiction.

On the one hand, the maximality of (B, A) in $\tilde{\tau} \cap \vec{S}_k$ implies that $(B \cup C \cup e, A \cap D) \notin \tilde{\tau}$. On the other hand, since $\{C, D\}$ is not a separation of G, but (C, D) is maximal in τ', the orientation (C, D) of $\{C, D\}$ cannot be forced by τ. Hence by construction of τ', we put $(C, D) \in \tau'$ because of $(C \cup e, D) \in \tilde{\tau}$. But $\{(B, A), (C \cup e, D), (A \cap D, B \cup C \cup e)\}$ is forbidden triple in G, so $(A \cap D, B \cup C \cup e) \notin \tilde{\tau}$. ■

Using Claim 2, we can now find a forbidden triple $\{(C_1, D_1), (C'_2, D'_2), (C'_3, D'_3)\}$ in τ as follows: As shown above, (C_1, D_1) is in τ and satisfies $(A, B) \leqslant (C_1, D_1)$. If (C_i, D_i) with $i \in \{2, 3\}$ is a separation of G, then it also is in τ, as τ extends to τ', and we set $(C'_i, D'_i) := (C_i, D_i)$. If it is not a separation of G, then $\{C_i, D_i\}$ must cross $\{A, B\}$, as shown above Claim 2, and Claim 2 yields $(C'_i, D'_i) := (B \cap C_i, A \cup D_i) \in \tau$. To see that $\{(C_1, D_1), (C'_2, D'_2), (C'_3, D'_3)\} \subseteq \tau$ is indeed a forbidden triple in G, note that

$$G[C_1] \cup G[C'_2] \cup G[C'_3] \supseteq G[C_1] \cup G[C_2 \cap B] \cup G[C_3 \cap B] \supseteq G[C_1] \cup G[C_2 \cap D_1] \cup G[C_3 \cap D_1] = G,$$

where the last equation holds because $e \in G[A] \subseteq G[C_1]$ and $\{(C_i, D_i) : i \in \{1, 2, 3\}\}$ is a forbidden triple in $G' = G - e$. This concludes the proof. □

Proof of Theorem 4.4.1. This follows immediately from Lemma 4.4.3 and Lemma 4.4.4. □

Our proof of Theorem 4.4.1 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 4.4.5. *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 $G - e$?*

4.5 Rainbow-cloud-decompositions in the absence of high-order tangles

Recall that Theorem 4.4.1 in Section 4.4 immediately yields Theorem 15 if the graph has a $(k+1)$-tangle. So from now on we work towards a proof of Theorem 15 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 [KNTW19]. We will later use that this decomposition exhibits a substructure of the graph, the 'rainbow', which is a long linear structure that is fairly independent of the rest of the graph and internally made up of similar enough parts such that deleting an edge in one of the parts does not change the overall structure of the graph. In particular, it will allow us to understand how the separations of the graph change after deleting such an edge and hence how to find a tangle of this smaller graph to which our given tangle extends.

We begin by building up to the definition of a 'rainbow-cloud decomposition'. Let G be a graph. First, a *linear decomposition*[3] of G of *length* $M \in \mathbb{N}$ of G is a family $\mathcal{W} = (W_0, W_1, \dots, W_M)$ of sets W_i of vertices of G such that

(L1) $\bigcup_{i=0}^{M} G[W_i] = G$,

(L2) if $0 \leqslant i \leqslant j \leqslant k \leqslant M$, then $W_i \cap W_k \subseteq W_j$, and

(L3) there is an integer ℓ such that $|W_{i-1} \cap W_i| = \ell$ for every $i \in [M]$, and

(L4) $W_{i-1} \neq W_{i-1} \cap W_i \neq W_i$ for all $i \in [M]$.

We call the sets W_i the *bags* and and the induced subgraphs $G[W_i]$ the *parts* of the linear decomposition $\mathcal{W}$. Note that the bags of a linear decomposition of length at least 1 are non-empty by (L4). The *adhesion sets* of a linear decomposition $\mathcal{W}$ are the sets $U_i := W_{i-1} \cap W_i$ for $i \in [M]$. The size of the adhesion sets U_i is the *adhesion* of $\mathcal{W}$. We emphasise that adhesion 0 is allowed. Whenever we introduce a linear decomposition as $\mathcal{W}$ without specifying its bags, then we will tacitly assume the bags to be denoted by $W_0, \dots, W_M$ and the adhesion sets by $U_1, \dots, U_M$.

Next we turn to the definition of 'rainbow-decompositions' which are special linear decompositions whose adhesion sets are minimal U_1–U_M separators of G as witnessed by respective families of disjoint paths. To make this formal, a *linkage* in a graph G is a set $\mathcal{P}$ of disjoint paths in G. If A and B are sets of vertices of G such that $\mathcal{P}$ consists of A–B paths, i.e. such paths that meet A precisely in one endvertex and B precisely in the other endvertex, then $\mathcal{P}$ forms an *A–B linkage*. A linear decomposition $\mathcal{W}$ of adhesion ℓ and length M is called a *rainbow-decomposition* of *adhesion* ℓ and *length* M if it has the following three properties:

(R1) There exists a U_i–U_{i+1} linkage of cardinality ℓ in $G[W_i]$ for every $i \in [M-1]$.

[3] Decompositions satisfying (L1) and (L2) are often known as *path-decompositions* (cf. [D17b, §12.6]). In [KNTW19] these are referred to as *linear decompositions*. In this chapter, linear decompositions will always not only satisfy (L1) and (L2) but also (L3) and (L4).

(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 4.2 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.

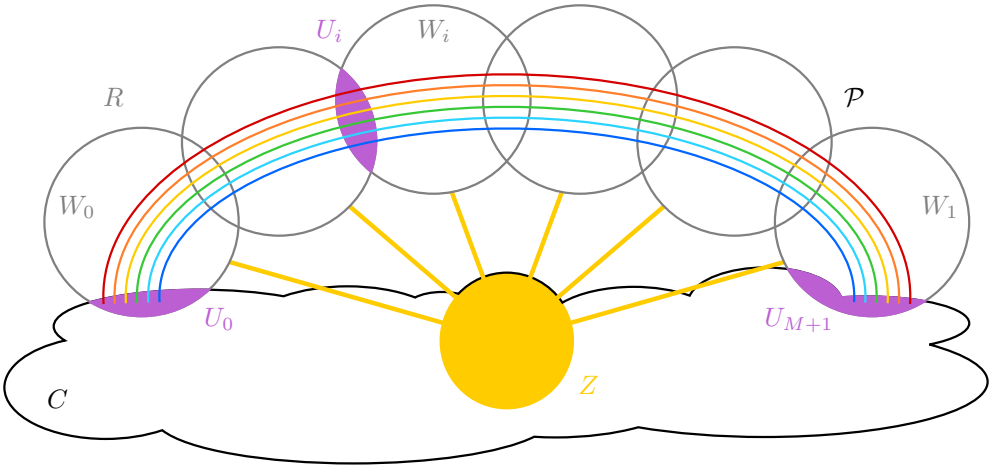

Figure 4.2: A rainbow-cloud decomposition $(R, \mathcal{W}, Z, C)$: the rainbow R together with its decomposition $\mathcal{W}$ is indicated in grey, the foundational linkage $\mathcal{P}$ of $\mathcal{W}$ is depicted in rainbow colours, and the cloud C is depicted in black. Further, indicated in yellow, 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 4.5.1. *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 4.5.1, 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 4.5.3 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. [D17b, 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 4.5.2. *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

$$(a_1, \dots, a_{i_1 - 1}), (a_{i_1 + 1}, \dots, a_{i_2 - 1}), \dots, (a_{i_{n'} + 1}, \dots, a_p) \tag{4.1}$$

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 (4.1), let us call it $(b_1, \dots, b_{p'})$, must have length at least $p' \geqslant n^{m-1}$, and we thus can apply the induction hypothesis to $(b_1 - 1, \dots, b_{p'} - 1)$ with integers $n, m - 1, p'$ to obtain a subsequence of $(b_1, \dots, b_{p'})$ which is as desired. □

A sequence $((A_i, B_i))_{i\in[p]}$ of (oriented) separations of a graph G is *strictly increasing* if $(A_i, B_i) < (A_j, B_j)$ for every two elements $i < j$ of $[p]$.

Lemma 4.5.3. *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. For each strictly increasing sequence T in $\vec{S}_m(G)$, we define an associated 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 4.5.2 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 4.5.2. 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

$$((C_{i_j}, D_{i_j})), (C_{i_j+1}, D_{i_j+1}), \ldots, (C_{i_{j+1}}, D_{i_{j+1}}))$$

of $((C_i, D_i))_{i\in[p']}$; for notational simplicity, we also denote this subsequence by $R = ((A_i, B_i))_{i\in[r]}$. Then we can obtain a new strictly increasing sequence $R' = ((A'_i, B'_i))_{i\in[r']}$ by removing duplicates, which will only appear consecutively, from the sequence

$$(A_1 \cap A, B_1 \cup B), \ldots, (A_q \cap A, B_q \cup B), (A, B), (A_1 \cup A, B_1 \cap B), \ldots, (A_r \cup A, B_r \cap B).$$

We now consider the strictly increasing sequence T' which is obtained from T by replacing its subsequence R with R'. As every element of the distributive lattice $\vec{U}(G)$ is uniquely determined by its infimum and supremum with any given other element in $\vec{U}(G)$, T' has length at least $q + 1 \geqslant p$; thus, $T' \in \mathcal{T}$. Moreover, the sequence of integers associated to T' is larger than the one associated with T, which contradicts the choice of T. Indeed by the choice of R, all separations (C_g, D_g) with $i_j \leqslant g \leqslant i_{j+1}$ have order at least ℓ. Hence, $n_i(T') \geqslant n_i(T)$ for every $i < \ell$ and moreover $n_{|A,B|}(T') \geqslant n_{|A,B|}(T) + 1$, since the new sequence T' additionally contains (A, B). □

To proceed to the proof of Theorem 4.5.1, 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. [D17b, 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 [EKT22b, Lemma 2.7] or [D18b; 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 4.5.4 (e.g. [D17b, 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 theorem to obtain a tree-decomposition of small width and then sort out a suitable linear decomposition by Lemma 4.5.3.

Lemma 4.5.5. *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 4.5.5. 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 4.5.3, 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. [D17b, 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, $(A_i, B_i) < (C, D) < (A_{i+1}, B_{i+1})$, which contradicts our choice of the (A_j, B_j) via Lemma 4.5.3.

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 4.5.5 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]:

[4](FL1) and (FL2) are (L7) and (L8) in [KNTW19], respectively.

(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 4.5.5.

Lemma 4.5.6 ([KNTW19, Lemma 3.5][5])**.** *For every two integers $M \geqslant 1$ and $\ell \geqslant 0$, there exists an integer $M_1 = M_1(\ell, M) \geqslant 1$ such that if a linear-decomposition $\mathcal{W} = (W_0, \dots, W_{M_1})$ 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).*

We remark that $M_1(\ell, M) = (M^{\binom{\ell}{2}+1} \cdot \binom{\ell}{2}!)^{\ell+1} \cdot \ell!$ suffices, as the proof of Lemma 4.5.6 in [KNTW19] shows.

We now turn to the proof of Theorem 4.5.1. 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 4.5.1. Suppose that G has no $(k+1)$-tangle. We then set $N(k, M) := N_1(k, M_1(k, M+2))$, where N_1 and M_1 are as in Lemma 4.5.5 and Lemma 4.5.6, respectively. Lemma 4.5.5 now 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 4.5.6, 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 $H_{\mathcal{P}'}$,

[5] More precisely, it follows from the proof of [KNTW19, Lemma 3.5], as (L6) is only assumed to guarantee that the resulting decomposition also satisfies (L6).

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. □

4.6 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 15. These tools will allow us to control the new low-order separations that arise when we delete an edge deep inside the rainbow.

4.6.1 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 4.6.1. *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 $A \smallsetminus B$ and $B \smallsetminus A$. Hence, at most $2k$ bags of $\mathcal{W}$ are not contained in either $A \smallsetminus B$ or $B \smallsetminus A$. □

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 4.6.2. *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 $D \subseteq G[A \smallsetminus B]$ 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 4.6.1 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. □

4.6.2 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 4.6.3. *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 4.6.4. *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 $2\ell + |Z|$.* □

4.6.3 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 $W_i \subseteq A \smallsetminus B$ and $W_j \subseteq B \smallsetminus A$. 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), $W_i \subseteq A \smallsetminus B$ and $W_j \subseteq B \smallsetminus A$. 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 4.3): More precisely, let $(R, \mathcal{W}, Z, C)$ be an RC-decomposition of a 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 $W_i \subseteq A \smallsetminus B$ and $W_j \subseteq B \smallsetminus A$, 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 4.6.5. *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 $h \in \{i_{A,B} + 1, \dots, j_{A,B}\}$.*

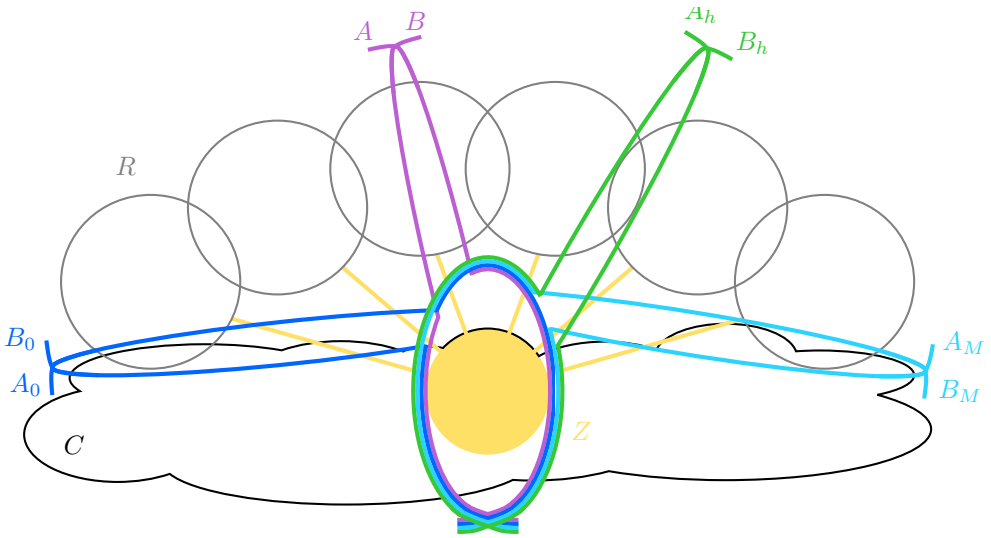

Figure 4.3: A rainbow-crossing separation (A, B) and the arising separations (A^0, B^0), (A^h, B^h) and (A^M, B^M).

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 4.6.4, 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 $\{A^h, B^h\}$:

$$\begin{aligned}
|A^h \cap B^h| &= |((A \cap V(C_{i,j}) \cup V(R_{i,h-1})) \cap (V(R_{h,j}) \cup (B \cap V(C_{i,j})))| \\
&\leqslant |(A \cap B) \cap V(C_{i,j})| + |A \cap (V(C_{i,j}) \cap V(R_{h,j}))| \\
&\quad + |B \cap (V(C_{i,j}) \cap V(R_{i,h-1}))| + |V(R_{i,h-1}) \cap V(R_{h,j})| \\
&\leqslant (k - \ell) + |A \cap W_j| + |B \cap W_i| + |U_h| = (k - \ell) + 0 + 0 + \ell = k,
\end{aligned}$$

where the penultimate equation holds because $W_i \subseteq A \setminus B$ and $W_j \subseteq B \setminus A$. This completes the proof. □

4.6.4 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 4.6.6. *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 $B \smallsetminus A$ 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, $W_h \subseteq B \smallsetminus A$ 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 $W_h \subseteq B \smallsetminus A$. 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 4.6.7. *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 4.6.8. *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 $W_i \subseteq A \smallsetminus B$ and $W_j \subseteq B \smallsetminus A$, 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 4.6.1, all but at most $2k$ bags of $\mathcal{W}$ are contained in either $A \smallsetminus B$ or $B \smallsetminus A$. 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 $A \smallsetminus B$ or $B \smallsetminus A$.

If one of W_h and W_s is contained in $A \smallsetminus B$ and the other one in $B \smallsetminus A$, 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 $W_h, W_s \subseteq A \smallsetminus B$. 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. □

4.7 Proof of Theorem 15 if G has no higher-order tangle

In this section, we complete the proof of Theorem 15. Based on the results in Sections 4.3 to 4.5, 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 15, reads as follows:

Theorem 4.7.1. *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 4.7.1. In the very end we combine the previous results to prove Theorem 15.

4.7.1 Living in the rainbow

In the proof of Theorem 4.7.1, 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 4.7.1, 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 4.7.2. *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 $B \smallsetminus A \subseteq V(R \smallsetminus C)$. 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 $Y \smallsetminus X \subseteq \bigcup_{t=r}^{s} W_t$ for integers $r \leqslant s$ with $s - r < 2k - 2$.*

Proof of the claim. 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 4.6.4.

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 $(V(C) \cup (V(R) \smallsetminus W_h) \cup U_h \cup U_{h+1}, W_h \cup Z) \in \tau$, which is the desired separation (X, Y) with $r := h =: s$.

Secondly, suppose that $2\ell + |Z| \geqslant k$, and let $(A, B) \in \tau$. We may assume that the separation $(A, B) \in \tau$ with $B \smallsetminus A \subseteq V(R) \smallsetminus V(C)$ is $\leqslant$-maximal in τ with that property; in particular, $G[B \smallsetminus A]$ is connected. Then $B \smallsetminus A$ cannot contain a bag of $\mathcal{W}$ by Lemma 4.6.6, since $B \smallsetminus A \subseteq V(R) \smallsetminus V(C)$ yields that $U_0, U_{M+1}, Z \subseteq V(C) \subseteq A$, and $\{A, B\}$ has order $< k \leqslant 2\ell + |Z|$. Thus, every bag which meets $B \smallsetminus A$ also meets the separator $A \cap B$. Thus, Lemma 4.6.1 ensures that at most $2k - 2$ bags of $\mathcal{W}$ meet $B \smallsetminus A$. Further, since $G[B \smallsetminus A] \subseteq R$ is connected and each bag W_i separates R, the bags which are met by $B \smallsetminus A$ are consecutive bags of $\mathcal{W}$. Thus, $(X, Y) := (A, B)$ is the desired separation, as witnessed by the respective indices r and s of the first and last bag of $\mathcal{W}$ which are met by $B \smallsetminus A$. ■

Let $(X, Y) \in \tau$ and $r \leqslant s$ with $s - r < 2k - 2$ be given by Claim 1. Now if $r > \frac{M}{2} - k$, then we set $i := 0$ and $j := r - 1$, and otherwise, if $r \leqslant \frac{M}{2} - k$, then we set $i := s + 1$ and $j := M$. We remark that in the latter case $s < \frac{M}{2} + k$. By Lemma 4.6.3, $(R, \mathcal{W}, Z, C)_{i,j}$ 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 of the claim. 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 4.6.3, $(R, \mathcal{W}, Z, C)_{i,j}$ 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. □

4.7.2 Deleting an edge and orienting the new separations

This subsection is dedicated to the proof of Theorem 4.7.1. We remark that this section is the only part of the proof of Theorem 15 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 4.7.1:

Setting 4.7.3. *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 4.7.4. *If we assume Setting 4.7.3 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 4.7.2 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 4.6.3, 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 4.6.4. 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 4.7.4 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 4.7.1, 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 4.4.4). 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 4.7.6 and Lemma 4.7.7, respectively). This then allows us to show that an even broader class of separations are forced by τ (Lemma 4.7.8). For all these subsequent three lemmas, we assume the following setting:

Setting 4.7.5. *Assume Setting 4.7.3 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 4.7.4 yields. In particular, $M \geqslant 8k$.*

Lemma 4.7.6. *Assume Setting 4.7.5. 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 4.6.5, $(A^{i+1}, B^{i+1}), (A^j, B^j)$ are again separations of G' of order at most k' which cross the rainbow clockwise. They are also separations of G by the choice of e: The bag $W_{M/2}$ is in $B^{i+1} \cap A^j$, since $i \leqslant 2k' \leqslant M/2 \leqslant M - 2k' \leqslant j$ because $M \geqslant 4k$ and $k' \leqslant k - 1$. As both $(A^{i+1}, B^{i+1}), (A^j, B^j)$ cross the rainbow, Z is contained in each of their separators. Thus, both endvertices of e are in any of the cases contained in $B^{i+1} \cap A^j$, as claimed.

Since τ does not live in the rainbow by the choice of $(R, \mathcal{W}, Z, C)$, the separation $\{A, B\}$ of G is oriented by τ monotonically over the rainbow. In particular, either $(A^{i+1}, B^{i+1}), (A^j, B^j) \in \tau$ or $(B^{i+1}, A^{i+1}), (B^j, A^j) \in \tau$. By definition of the (A^h, B^h) and the choice of i, j, we have $(A^{i+1}, B^{i+1}) \leqslant (A, B) \leqslant (A^j, B^j)$ and $V(R_{2k'+1,M-2k'-1}) \subseteq B^{i+1}, A^j$. So either (A^j, B^j) or (B^{i+1}, A^{i+1}) can be chosen as the desired separation $(X, Y) \in \tau$, as $k' < k$. □

Lemma 4.7.7. *Assume Setting 4.7.5. 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 4.6.2 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 4.6.7 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 4.6.7. 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 4.7.8. *Assume Setting 4.7.5. 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 4.6.2 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 of $(R, \mathcal{W}, Z, C)$, it follows that $(V(C_A) \cup V(C_B)) \cap V(C) \neq \emptyset$. As either both or none of C_A and C_B meets C by assumption, both C_A and C_B meet C.

All in all, both C_A and C_B meet C and $W_{M/2}$. Hence, Lemma 4.6.2 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 4.6.8 ensures that $\{A, B\}$ either crosses or slices the rainbow in G'. Thus, the previous Lemmas 4.7.6 and 4.7.7 guarantee that τ forces an orientation of $\{A, B\}$. □

With these tools at hand, we are now ready to prove Theorem 4.7.1:

Proof of Theorem 4.7.1. We may assume Setting 4.7.5, as Setting 4.7.3 with $M_0 \geqslant 18k$ is precisely the premise of Theorem 4.7.1. 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 4.7.8 ensures that precisely one of C_A and C_B meet C for a separation $\{A, B\}$ in case (2). It is immediate form (1)that τ' contains τ; so τ extends to τ'. It remains to show that τ' is indeed a k-tangle in G'.

Suppose for a contradiction that τ' is not a tangle in G', i.e. the orientation τ' of $S_k(G')$ contains a forbidden triple $\{(A'_i, B'_i) : i \in [3]\} \in \mathcal{T}(G')$. By the definition of $\mathcal{T}(G')$, we

may assume without loss of generality that all the (A'_i, B'_i) are $\leqslant$-maximal in τ'. In the remainder of the proof, we will construct from the forbidden triple $\{(A'_i, B'_i) : i \in [3]\} \subseteq \tau'$ a forbidden triple $\{(A_i, B_i) : i \in [3]\} \subseteq \tau$, which then contradicts that τ is a tangle in G.

To this end, we will make use of the RC-decomposition $(R, \mathcal{W}, Z, C)$. We first show that one of the (A'_i, B'_i) contains $V(R_{2k-1,M-2k+1}) \cup Z$ in its respective small side A'_i; in particular, the edge e is contained in $G[A'_i]$ and thus it is a separation of G. We then show for the other two separations (A'_i, B'_i) in the forbidden triple that either they are separations of G, too, or the component $C_{A'_i}$ of $G[A'_i \setminus B'_i]$ containing an endvertex of e in fact contains also $V(R_{2k-1,M-2k+1})$. Moving the $C_{A'_i}$ to the respective big side B'_i if necessary, these three separations of G will then yield a forbidden triple in τ.

So we first show the following:

Claim 1. *If $\{(A'_i, B'_i) : i \in [3]\} \in \mathcal{T}(G')$ is contained in τ' where each (A'_i, B'_i) is $\leqslant$-maximal in τ', then some (A'_j, B'_j) is also a separation of G with $V(R_{2k-1,M-2k+1}) \cup Z \subseteq A'_j$. In particular, $(A'_j, B'_j) \in \tau$.*

Proof of the claim. For every (A'_i, B'_i), at most $2k-2$ bags of $\mathcal{W}$ are contained neither in $A'_i \setminus B'_i$ nor in $B'_i \setminus A'_i$ by Lemma 4.6.1. Thus, all but at most $6k-6$ bags are contained in one of the strict sides of every (A'_i, B'_i). As $\{(A'_i, B'_i) : i \in [3]\}$ is a forbidden triple in G', we in particular have that $R[A'_1]$, $R[A'_2]$ and $R[A'_3]$ together cover the rainbow R in G'. Thus, no such bag may be contained in every big side B'_i. Since $M \geqslant 6k-6$, there thus must be some bag W_h of $\mathcal{W}$ that is contained in the strict small side $A'_j \setminus B'_j$ of some (A'_j, B'_j); by renaming the (A'_i, B'_i), we may assume that $j = 1$. Note that as $W_h \subseteq A'_1 \setminus B'_1$ and $Z \subseteq N_{G'}(W_h)$ by (RC4), we have $Z \subseteq A'_1$. We claim that (A'_1, B'_1) is as desired, i.e. it is also a separation of G and $V(R_{2k-1,M-2k+1}) \subseteq A'_1$.

Assume that $B'_1 \setminus A'_1$ also contains some bag W_s. By Lemma 4.6.8, (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 4.7.6 and 4.7.7. 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 4.7.6 and 4.7.7 yield $V(R_{2k-1,M-2k+1}) \subseteq A'_1$, as desired.

So now assume that no bag of $\mathcal{W}$ is contained in $B'_1 \setminus A'_1$. By Lemma 4.6.1, the strict side $A'_1 \setminus B'_1$ then contains all but at most $2k-2$ many bags of $\mathcal{W}$. Suppose for a contradiction that $\{A'_1, B'_1\}$ was not already a separation of G. Since (A'_1, B'_1) is $\leqslant$-maximal in τ', the tangle τ 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 4.6.2 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 4.6.2 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 of the claim. 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 4.6.2 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 4.6.1. 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 4.6.8. Thus, τ forces an orientation of $\{A, B\}$ by Lemmas 4.7.6 and 4.7.7, 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 4.6.1 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 included 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 15. Choose $M(k)$ to be $N(k, 18k)$ as in Theorem 4.5.1. 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 4.3.1 and 4.3.2. If there exists a $(k+1)$-tangle in G, then we are done by Theorem 4.4.1. 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 4.5.1. If G has a vertex of degree at most 2, then we are done by Lemmas 4.3.4 and 4.3.5. Thus, Theorem 4.7.1 concludes the proof. □

We remark that one may calculate that $M(k) \in O(3^{k^{k^5}})$.

4.8 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 14. Next, we deduce from Theorem 14 our reduction of Problem 4.1.1 to small graphs, Theorem 11, and then derive Corollaries 12 and 13 from it. Finally, we present a further application of Theorem 14 in Corollary 4.8.4, 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 4. *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 15. Suppose that $G_0, \ldots, G_{i-1}$ and $\tau_0, \ldots, \tau_{i-1}$ are already defined. If G_{i-1} is disconnected, then we apply Lemma 4.3.3 to obtain G_i and τ_i. If $k \leqslant 2$, then we apply Lemma 4.3.1 or Lemma 4.3.2. If $k \geqslant 3$, but G_{i-1} has a vertex of degree $\leqslant 2$, then we apply Lemma 4.3.4 or Lemma 4.3.5. 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 15, if G_{i-1} has at least $M(k)$ edges. Otherwise, we set $m := i - 1$, completing the proof. □

4.8.1 Application I: sets and functions inducing tangles

In this section, we address Problem 4.1.1.

Problem 1.1. *Is every tangle in a graph G induced by some set $X \subseteq V(G)$?*

We will use Theorem 4 to prove Theorem 11, our reduction of Problem 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 11 reads as follows:

Theorem 1. *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 1 via Theorem 4, 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 4.8.1. *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 4.8.2. *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 4.8.3. *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 1. Let τ_m be the k-tangle in the graph G_m with less than $M(k)$ edges as described in Theorem 4. Let w' be a weight function on $V(G')$ which induces the tangle τ_m. Iteratively applying Lemmas 4.8.1 to 4.8.3 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 1, we now deduce Corollary 2 and Corollary 3:

Corollary 2. *For $k \geqslant 1$, there exists $M = M(k) \in O(3^{k^{k^5}})$ such that Problem 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 1. □

Corollary 3. *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 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 11.

Proof. We enumerate all the finitely many non-isomorphic connected graphs $G_1, \ldots, G_m$ with fewer than M edges. As every finite graph has at most finitely many tangles, there are only finitely many k-tangles τ in any such G_i. Elbracht, Kneip and Teegen showed with Theorem 4.1.2 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 1 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 1 by choosing the weight functions as indicator functions of the inducing sets given by the assumed positive answer to Problem 1.1. □

We remark that the proof of Corollary 3 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 1.

4.8.2 Application II: subgraphs witnessing a tangle

In this section, we demonstrate another application of our inductive proof method Theorem 4 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][6] proved that every k-tangle is witnessed by a set of edges whose size can be bounded in k. However, their bound is defined recursively and yields a power tower of height $k-1$. By Theorem 4, we obtain a new bound which is significantly better for sufficiently large k:

Corollary 4.8.4. *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 4.8.4. Let $M(k)$ be given by Theorem 4, and set $M'(k) := 2M(k)$. Let τ be a k-tangle in some graph G. Theorem 14 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 4.

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 4 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. □

[6] We remark that triple covers in their paper are precisely the witnessing sets here. In fact, they proved a more general result about tangles on bipartitions in a more general setting. However, every k-tangle τ in G induces a tangle τ' on the set of bipartitions of the edge set $E(G)$ of order $< k$: let $(C, D) \in \tau'$ if and only if there exists a separation $(A, B) \in \tau$ with $C \subseteq E(G[A])$ and $D \subseteq E(G[B])$. Then their result yields the described conclusion.

CHAPTER 5

Point sets and functions inducing tangles of set separations

This chapter is based on [DEJ24] and joint work with Reinhard Diestel and Christian Elbracht.

Tangles, as introduced by Robertson and Seymour, were designed as an indirect way of capturing clusters in graphs and matroids. They have since been shown to capture clusters in much broader discrete structures too. But not all tangles are induced by a set of points, let alone a cluster. We characterise those that are: the tangles that are induced by a subset of or function on the set of data points whose connectivity structure they are meant to capture.

We offer two such characterisations. The first is in terms of how many small sides of a tangle's separations it takes to cover the ground set. The second uses a new notion of duality for oriented set separations that becomes possible if these are no longer required to be separations of graph or matroids.

5.1 Introduction

Tangles were introduced by Robertson and Seymour as a tool in their graph minors project [RS91]. They provided a novel, indirect, way to capture highly cohesive substructures, or 'clusters', in graphs. The idea is that since clusters cannot be divided into significantly large parts by graph separations of low order, any given cluster implicitly orients every low-order separation towards its 'big' side, the side that contains most of the cluster. It turned out that this induced orientation of all the low-order separations collectively contains all the information needed to prove fundamental theorems about the cluster structure of a graph, which has made tangles a powerful tool in the connectivity theory of graphs.

Over the last decade, the theory of tangles has been significantly generalised to other discrete structures [A24a; DEW19; EK21; EKT21; EKT22a; EKT22b]. These include matroids [GGRW06], but also bespoke structures that come with concrete clustering applications [D24; KEFKRTL23]. This has been made possible by re-casting the notion and theory of tangles in terms of a purely algebraic framework of 'separation systems' [D18a; DHL19; DO21], which encompass the notions of separation from all these various different contexts. Although these separation systems are very general, the central tangle theorems are still valid in this framework.

In this chapter we are concerned with only the most basic type of separation systems, those of sets. A *separation* of a set V is an unordered pair $s = \{A, B\}$ of subsets of V such that $A \cup B = V$. It has two *orientations*: the ordered pair (A, B), which we think of as pointing towards B, and its *inverse* (B, A). We usually denote the two orientations of a separation s by arrows: one of them is denoted as $\vec{s}$, the other as $\overleftarrow{s}$, but it does not matter which is which.

Now consider a particular set S of separations of a set V, and a subset $X \subseteq V$ that is a 'cluster' in the sense that the separations in S cannot divide it evenly: let us assume that for every $\{A, B\} \in S$ more than two thirds of X lies in $A \setminus B$ or in $B \setminus A$. If X has more elements in B than in A, we think of X as orienting this separation towards B; in our notation it orients it 'as (A, B)'.

Any orientation τ of (all the separations in) S induced by a cluster $X \subseteq V$ in this way has the following property, which no longer refers to X: whenever τ orients three separations $\{A_i, B_i\} \in S$ $(i = 1, 2, 3)$ towards B_i, their 'small sides' A_i cannot cover V, because each contains less than a third of X. This is the most basic definition of a *tangle* of S: any orientation of S such that no three small sides cover V. Note that the meaning of 'small' here is intrinsic to τ: the side of a separation towards which τ orients it is now called 'big', its other side 'small'. No reference to a 'cluster' X is made in this definition of a tangle.

This abstract definition of a tangle has made it possible to investigate clusters in a graph or data set without referring to them directly in the usual concrete way, as sets of vertices or data points. In particular, one can investigate the relative structure of clusters without even having found them in this concrete sense – a sense which, moreover, may be inadequate given the fuzzy nature of many real-world clusters, which often do not enable us to decide easily which data points belong to a cluster and which do not.

However, tangles are not equivalent to clusters but more general: while every cluster, in our earlier informal sense, gives rise to a tangle, not every tangle is induced by such a cluster. We shall see an example in a moment. Tangles that do not come from clusters can still be interesting; the *text tangles* in [D24; DEK] are a typical example in the context of set separations.

To be a little more formal, let us say that a set $X \subseteq V$ of 'points' *induces* an orientation τ of S if, for every $(A, B) \in \tau$, there are more elements of X in B than in A. It is an open question whether every tangle of a graph is induced by some set of its vertices.[1] Tangles of more general set separations, however, need not be induced by a set of points. Let us construct a simple example.

The basic idea of our construction is that we start with S and an 'orientation' $\tau = \{\vec{s} \mid s \in S\}$ of S as just a collection of names. Every $\vec{s}$ will eventually become a pair (A, B) of subsets of some set V to be constructed; a pair of subsets that form a separation of V. In order to flesh out our notational shells for $\tau = \{\vec{s} \mid s \in S\}$ and each $\vec{s} = (A, B)$ in a way that makes τ into a tangle of S not induced by any subset of V, we build V element by element, immediately assigning every $v \in V$ to either A or B for

[1] It was shown in [EKT20] that graph tangles are induced by sets of vertices with weights assigned to them; we then say that these weight functions induce those tangles; see later.

every $\vec{s} = (A, B) \in \tau$. When we are done, we shall verify that τ is indeed a tangle of S not induced by any set in V.

To implement this plan, we create for every 3-set $\sigma \subseteq \tau$ one element v_σ for V so that these v_σ are distinct for different σ, and we let V be the set of all these v_σ. For each $(A, B) \in \sigma$ we put v_σ in B, and for each $(A, B) \in \tau \setminus \sigma$ we put v_σ in A. Then for every $\vec{s} = (A, B) \in \tau$ we have $B = \{ v_\sigma \mid \vec{s} \in \sigma \}$ and $A = V \setminus B$. Now if $|S| = m$, say, then $|V| = \binom{m}{3}$, and for every $(A, B) \in \tau$ we have $|B| = \binom{m-1}{2}$. Let us choose S so that $m \geqslant 6$.

By construction, $\tau = \{\vec{s} \mid s \in S\}$ is a tangle of S: the big sides of any three $\vec{s} \in \tau$ have an element in common. Conversely,

$$\textit{every } v \in V \textit{ lies on the big side of exactly three elements of } \tau. \tag{$*$}$$

A simple double count now shows that τ is not induced by any set $X \subseteq V$. Indeed, suppose it is and let

$$d := \sum_{(A,B)\in\tau} |X \cap B|.$$

Then $|X \cap B| > |X|/2$ for every $(A, B) \in \tau$, because X induces τ, and hence $d > m\,|X|/2$. On the other hand, by $(*)$, each $x \in X$ lies in B for exactly three $(A, B) \in \tau$, so d counts it three times: $d = 3\,|X|$. Putting these together we obtain $3\,|X| > m\,|X|/2$. This implies $m < 6$, contrary to our assumption.

Is it possible to distil from this example some property of τ that identifies *all* the tangles that are not induced by any set of 'points', elements of their ground set? We grappled with this question for quite a while, until we found the following solution.

Given an integer k, we say that an orientation τ of S is *k-resilient* if it takes more than k elements of τ to obtain V as the union of their small sides. Every tangle, by definition, is 3-resilient.

Our earlier example of a tangle τ not induced by a set in V is not 4-resilient. In fact, given any four distinct separations in τ, by $(*)$ every $v \in V$ lies on the small side of at least one of them, so the four small sides have union V. At the other extreme, every *principal* tangle of a set of bipartitions of a set V, one consisting of all bipartitions (A, B) of V whose big side B contains some fixed element $x \in V$, is *infinitely resilient* in that it is k-resilient for every $k \in \mathbb{N}$. Note that $\{x\}$ induces this tangle. In Section 5.3 we shall see that the unique 5-tangle of the $(n \times n)$-grid, which is induced by its entire vertex set, is $\Omega(n^2)$-resilient.

These examples seem to suggest that tangles of set separations that are k-resilient for large k are more likely to be induced by subsets of their ground set. We can indeed prove such a fact, with an interesting additional twist: 'large' has to be measured not in terms of $|V|$ or $|S|$, but relative to the number of maximal elements of the tangle in the usual partial order of oriented separations (see Section 5.2). This dependence on the number of maximal elements in a tangle is quite natural: in a k-resilient tangle with at most k maximal elements, the intersection of all their big sides is non-empty, and it clearly induces this tangle.

Let us say that a function $w\colon V \to \mathbb{R}_{\geqslant 0}$, which we may think of as placing weights on the elements of V, *orients* a separation $s = \{A, B\}$ of V *as* $\vec{s} = (A, B)$ if $w(A) < w(B)$, where

$w(A) = \sum_{v \in A} w(v)$ and likewise for B. More generally, w *orients* a set S of separations of V *as* $\{\vec{s} \mid s \in S\}$ if it orients every $s \in S$ as $\vec{s}$.

Conversely, if τ is a given orientation of S and $w\colon V \to \mathbb{R}_{\geqslant 0}$ orients S as τ, we say that w *induces* τ. If w induces τ and takes values in $\{0, 1\}$, we likewise say that $X = w^{-1}(1) \subseteq V$ *induces* τ; note that this agrees with our earlier informal notion of tangle-inducing subsets of V.

Our earlier observation that k-resilient tangles with at most k maximal elements are induced by point sets can be strengthened for inducing functions:

Theorem 5.1.1. *A tangle with m maximal elements is induced by some function if it is k-resilient for some $k > \frac{m}{2}$.*

We shall see that, as a general bound for all tangles, this is best possible.

For our proof of Theorem 5.1.1 we introduce the notion of being 'locally induced', which generalises the idea of resilience. We show in Theorem 5.3.5 that an orientation τ of a set S of separations is induced by a function if and only if there exist suitable parameters k and ℓ such that τ is 'k-locally ℓ-induced'. In particular, there are such suitable parameters k and ℓ if τ is highly resilient compared with its number of maximal elements, our Theorem 5.1.1 above.

For our second characterisation of tangles induced by functions we exploit the recent notion of *duality* of sets of separations. Duality between oriented set separations naturally arises in applications of tangle theory. For example, consider an online shop with a set V of items on sale and a history P of purchases made last year [D19; D24]. In this setting, two different sets of separations occur. Every purchase in P induces a bipartition of V into the items bought versus those not bought. Equally, every item in V defines a bipartition of P into those purchases that included it versus those that did not. The tangles of these two sets of separations can be shown to interact [DEET23], and they will help us to obtain a second characterisation of tangles induced by functions, Theorem 5.4.1. This will also imply Theorem 5.1.1.

In our third contribution in this chapter we show that some tangles are induced by point sets if the separations they orient are endowed with an *order function*, a function that assigns an integer $|s| =: |A, B| \geqslant 0$ to every separation $s = \{A, B\} \in U$ where $U = U(V)$ is the set of all separations of a set V (see Section 5.2 for the precise definitions). Given such an order function on U and an integer $k \geqslant 1$, the *k-tangles* in U are the tangles of its subset $S_k := \{\, s \in U : |s| < k \,\}$.

Elbracht, Kneip, and Teegen [EKT20] showed that the k-tangles in U are induced by functions on V, though not necessarily by subsets of V, when their order is defined as $|A, B| := |A \cap B|$. In Section 5.5 we strengthen their result for k-tangles in U which extend to $2k$-tangles in U, by showing that such k-tangles are even induced by sets. It would be interesting to know whether similar results hold for other submodular (see Section 5.2) order functions on U than the above.

5.2 Preliminaries

This section collects together the definitions we need in this chapter. While we shall work only with separations of sets as considered in the introduction, all definitions given here fit into the more general framework of 'abstract separation systems' [D18a]. From this framework we shall borrow some notations which we will also introduce in what follows.

In this chapter we consider separations of finite sets V, this finiteness assumption on V will not be mentioned explicitly in the remainder of this chapter. Throughout we will consider various finite sets V; if V is not specified explicitly, then V is any arbitrary finite set.

For definitions around graphs we refer the reader to [D17b]. For every $k \in \mathbb{N}$ we write $[k] := \{1, \dots, k\}$, we call a set with k elements a *k-set*, and we denote the set of all k-element subsets of a set X as $X^{(k)}$.

5.2.1 Separations of sets

An *(unoriented) separation* of V is an unordered pair $\{A, B\}$ of subsets A and B of V, its *sides*, such that $A \cup B = V$. The two *orientations* of $\{A, B\}$ are the ordered pairs (A, B) and (B, A) which is the *inverse* of (A, B).[2] We write $U(V)$ for the set of all unoriented separations of V.

Every ordered pair (A, B) of subsets of V with $A \cup B = V$ is an *oriented separation* of V. Its *underlying* unoriented separation is $\{A, B\}$, and (B, A) is its *inverse*. Given an oriented separation (A, B) of V, we refer to A as its *small side* and call B its *big side*. We shall informally use the term 'separation' also as a short term for oriented separations, but we will only do so if the meaning is unambiguous.

As indicated in the introduction, we fix the following notational conventions for separations for better readability: unoriented separations will be denoted as lowercase letters, such as s. Given an unoriented separation s of a set, we denote its two orientations as $\vec{s}$ and $\overleftarrow{s}$. There is no default orientation: once we have called one of the two orientations $\vec{s}$, the other one will be $\overleftarrow{s}$, and vice-versa. Oriented separations will be denoted as lowercase letters with a forward or backward arrow on top, such as $\vec{s}$ and $\overleftarrow{s}$. Given an oriented separation $\vec{s}$ of V, its underlying unoriented separation is denoted as s, and its inverse as $\overleftarrow{s}$.

We define a partial order $\leqslant$ on the set of oriented separations of V as follows: for two oriented separations (A, B) and (C, D) of V, we let $(A, B) \leqslant (C, D)$ if $A \subseteq C$ and $B \supseteq D$; we write $(A, B) < (C, D)$ if and only if $(A, B) \leqslant (C, D)$ and $(A, B) \neq (C, D)$. With this definition we in particular have

$$(A, B) \leqslant (C, D) \iff (B, A) \geqslant (D, C).$$

The *maximal elements* of a set σ of oriented separations of V are always those separations in σ which are maximal with respect to this partial order $\leqslant$.

Given two oriented separations $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$ of V, their supremum $\vec{r} \vee \vec{s}$ in $U(V)$ with respect to the partial order $\leqslant$ is the oriented separation $(A \cup C, B \cap D)$,

[2] Note that (A, B) and (B, A) coincide if and only if $A = B = V$. In particular, the oriented separation (V, V) equals its own inverse.

and their infimum $\vec{r} \wedge \vec{s}$ in $U(V)$ is $(A \cap C, B \cup D)$. Note that the supremum and the infimum satisfy DeMorgan's law in that $\overleftarrow{r} \vee \overleftarrow{s}$ is the inverse of $\vec{r} \wedge \vec{s}$ for every two oriented separations $\vec{r}$ and $\vec{s}$ of V.

A set U of unoriented separations of V is a *universe of separations of* V if the set $\vec{U} := \{\vec{s}, \overleftarrow{s} \mid s \in U\}$ of all orientations of separations in U is closed under taking suprema and infima in $U(V)$. Note that $U(V)$ itself is a universe of separations of V, by definition.

An oriented separation $\vec{s} = (A, B)$ of V is *small* if $B = V$, and *co-small* if $A = V$; thus, $\vec{s}$ is small if and only if its inverse $\overleftarrow{s}$ is co-small. Note that $\vec{s}$ is small if and only if $\vec{s} \leqslant \overleftarrow{s}$. If $A = V = B$ or, equivalently, if $\vec{s} = \overleftarrow{s}$, then both $\vec{s}$ and s are called *degenerate*; otherwise, $\vec{s}$ and s are *non-degenerate*. A set σ of oriented non-degenerate separations of V is a *star* if $\vec{r} \leqslant \overleftarrow{s}$ for every two distinct $\vec{r}, \vec{s} \in \sigma$.

Given a universe U of separations of V, an *order function* $|\cdot|$ on U assigns to each $s \in U$ its *order* $|s| \in \mathbb{Z}_{\geqslant 0}$; the order of an orientation $\vec{s}$ of a separation $s \in U$ is defined to be the order of s. Given $k \in \mathbb{N}$ we write S_k for the set of all separations in U of order less than k.

An order function on U is *submodular* if, for every two $r, s \in U$ and arbitrary orientations $\vec{r}$ of r and $\vec{s}$ of s, we have $|\vec{r} \vee \vec{s}| + |\vec{r} \wedge \vec{s}| \leqslant |\vec{r}| + |\vec{s}|$. Unless explicitly specified otherwise, we consider the *(standard) order* $|A, B|$ of a separation $\{A, B\}$ of V as the cardinality of its *separator* $A \cap B$; this order function can easily be seen to be submodular.

There are two special classes of separations which we will consider in this chapter: a *bipartition* of V is a separation of V whose sides are disjoint. We denote the set of all bipartitions of V by $U_{bip}(V)$; note that $U_{bip}(V)$ is again a universe of separations of V. Since all bipartitions have order 0 with respect to our standard order function, we usually consider other order functions on $U_{bip}(V)$ (see [D19] for various examples).

Another example of separations arises in graphs: a *separation* $\{A, B\}$ of a graph $G = (V, E)$ is a separation of its vertex set V such that G has no edges between $A \setminus B$ and $B \setminus A$. The set of all separations of G then forms a universe of separations of V.

5.2.2 Orientations

Let S be a set of unoriented separations of a set V. Assigning to every $s \in S$ either $\vec{s}$ or $\overleftarrow{s}$ is called *orienting* S (or the $s \in S$). So an *orientation* of S is a set τ of orientations of separations in S satisfying $|\tau \cap \{\vec{s}, \overleftarrow{s}\}| = 1$ for every $s \in S$. Given an orientation τ' of a subset $S' \subseteq S$, we say that τ' *extends* to an orientation τ of S if $\tau' \subseteq \tau$.

An orientation τ of S is *consistent* if there exist no two $\vec{r}, \vec{s} \in \tau$ with $\overleftarrow{r} < \vec{s}$. If τ is consistent and for every two distinct $\vec{s}, \vec{t} \in \tau$, we have $(\overleftarrow{s} \wedge \vec{t}) \notin \tau$, then τ is a *profile* of S. A profile of S is *regular* if it does not contain any co-small separation. Given a universe U of separations of V, an order function on U, and $k \in \mathbb{N}$, we call a profile of the corresponding S_k a *k-profile in U*.

Writing $\vec{S} := \{\vec{s}, \overleftarrow{s} \mid s \in S\}$ for the set of all orientations of separations in S, let $\mathcal{F} \subseteq 2^{\vec{S}}$. An orientation τ of S is an *$\mathcal{F}$-tangle* of S if τ is consistent and $\sigma \not\subseteq \tau$ for every $\sigma \in \mathcal{F}$. Let $\mathcal{T}$ be the set consisting of all sets $\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\}$ of (not necessarily distinct) oriented separations in $\vec{S}$ with $A_1 \cup A_2 \cup A_3 = V$, i.e. the supremum of the (A_i, B_i) is co-small. The $\mathcal{T}$-tangles of S are called *(abstract) tangles* of S; they are examples of regular profiles of S [DEW19]. Given an order function on a universe U of separations

of V and $k \in \mathbb{N}$, we call a tangle of $S_k \subseteq U$ a *k-tangle in U*.

In the case that all separations in S are even bipartitions of V, we will also consider $\mathcal{F}^\ell$-tangles where $\ell \in \mathbb{R}_{\geqslant 0}$.[3] Here, $\mathcal{F}^\ell \subseteq 2^{\vec{S}}$ consists of all sets $\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\}$ of (not necessarily distinct) oriented bipartitions in $\vec{S}$ with $|B_1 \cap B_2 \cap B_3| < \ell$. Note that $\mathcal{T} = \mathcal{F}^1$ here since S consists of bipartitions of V.

5.2.3 Point sets and functions inducing orientations

A *weight function* on a set V is a map w from V to $\mathbb{R}_{\geqslant 0}$. For a subset $Z \subseteq V$ we write $w(Z) = \sum_{v \in Z} w(v)$. If there exists $v \in V$ with $w(v) > 0$, then w is *non-zero*. If w takes values in $\{0, 1\}$ only, then it can equivalently be formulated as an indicator function of the set $X = X_w = w^{-1}(1)$ in that $w(Z) = |X \cap Z|$ for every $Z \subseteq V$; we shall use this equivalence freely throughout.

For any weight function w on V and any separation $\{A, B\}$ of V, we have $w(B) - w(A) = w(B \smallsetminus A) - w(A \smallsetminus B)$, a fact we shall use freely throughout. We say that w *orients* a separation $\{A, B\}$ of V *as* (A, B) if $w(A) < w(B)$. More generally, w *orients* a set S of separations of V *as* an orientation τ of S if $w(A) < w(B)$ for all $(A, B) \in \tau$.

Conversely, let S be a set of separations of a set V, and let τ be an orientation of S. If a weight function w on V orients S as τ, then we say that w *induces* τ and all its elements. If w induces τ and takes values in $\{0, 1\}$, then we say that the set $X = w^{-1}(1)$ *induces* τ.

Let us note some basic observations about orientations induced by functions. First, let w be a weight function on V and let $\lambda > 0$ be a scalar. If we *scale w by λ*, i.e. if we consider the weight function $v \mapsto \lambda w(v)$ on V, then this scaled weight function agrees with w on the sign of $w(B) - w(A)$ for any separation $\{A, B\}$ of V. In particular, if an orientation τ of a set S of separations of V is induced by a function w, then for any given $K > 0$ there exists a function w_K inducing τ with $w_K(B) - w_K(A) \geqslant K$ for all separations $(A, B) \in \tau$. This is because w_K can be chosen as an appropriate scaling of w, i.e. by a factor $\lambda \geqslant K/(\min_{(A,B)\in\tau}(w(B) - w(A)))$.

This fact directly implies that if an orientation τ of a set S of separations of V is induced by a function w, then there also exists a function inducing τ which takes values in $\mathbb{Z}_{\geqslant 0}$ instead of $\mathbb{R}_{\geqslant 0}$. Indeed, there exists $\varepsilon > 0$ such that $w(B) - w(A) \geqslant \varepsilon$ for all $(A, B) \in \tau$. Since $\mathbb{Q}$ is dense in $\mathbb{R}$, we can replace for every $v \in V$ the scalar $w(v) \in \mathbb{R}_{\geqslant 0}$ with a rational number $w'(v) \in \mathbb{Q}_{\geqslant 0}$ which satisfies $|w(v) - w'(v)| < \varepsilon/|V|$. By construction, the resulting weight function w' on V still induces τ. Now we can scale w' by an appropriate $\lambda \in \mathbb{N}$ to obtain the desired function which takes values in $\mathbb{Z}_{\geqslant 0}$ and induces τ.

For the final observation in this section, let S be a set of separations of a set V, and let τ be an orientation of S. Then a weight function on V induces τ as soon as it induces the maximal elements of τ. We include a proof of this observation from [EKT20] for the reader's convenience.

Lemma 5.2.1. *Let w be a weight function on a set V, and let (A, B) and (C, D) be separations of V with $(A, B) \leqslant (C, D)$. Then $w(B) - w(A) \geqslant w(D) - w(C)$. In particular, w induces (A, B) if it induces (C, D).*

[3] We slightly deviate here from [DEW19] in that our $\mathcal{F}^\ell$ correspond to their $\mathcal{F}_3^\ell$ and in that we consider $\ell \in \mathbb{R}_{\geqslant 0}$ instead of $\ell \in \mathbb{N}$.

Proof. Since $(A, B) \leqslant (C, D)$, we have $A \subseteq C$ and $B \supseteq D$. So $w(A) \leqslant w(C)$ and $w(B) \geqslant w(D)$ as w is a weight function on V. This directly implies that $w(B)-w(A) \geqslant w(D)-w(C)$. The 'in particular'-part then follows immediately. □

5.3 Resilience and locally induced orientations

In this section we use the novel notion of *resilience* to prove a sufficient criterion for an orientation of a set S of separations to be induced by a function. After that, we further generalise the concept of resilience towards the notion of being 'k-locally ℓ-induced' which allows us to give a characterisation of those orientations of S which are induced by functions on the ground set V. We begin this section by giving all the definitions around the concept of resilience.

Let S be a set of separations of a set V, and let $k \in \mathbb{N}$. An orientation τ of S is *k-resilient* if no set of at most k separations in τ has a co-small supremum. So τ is k-resilient if and only if for all sets $\sigma \subseteq \tau$ of size at most k, we have that $\bigcup_{(A,B)\in\sigma} A \neq V$. If S consists only of bipartitions of V, then this is equivalent to $\bigcap_{(A,B)\in\sigma} B \neq \emptyset$ for all sets $\sigma \subseteq \tau$ of size at most k since $(V, \emptyset)$ is the only co-small bipartition of V. Note that in order to determine whether τ is k-resilient, it is always enough to consider sets σ of maximal elements of τ.

If τ is k-resilient, then τ is also k'-resilient for every $k' < k$. We call τ *infinitely resilient* if τ is k-resilient for all $k \in \mathbb{N}$. The *resilience* of τ is the maximal $k \in \mathbb{N}$ such that τ is k-resilient if such k exists, 0 if τ is not k-resilient for any $k \in \mathbb{N}$, and ∞ otherwise.

In addition to the examples on resilience given in the introduction, let us here illustrate the concept once more with a less extreme example. Consider the $(n\times n)$-grid for some $n \geqslant 5$, and let S be the set of all separations of this graph which have order at most 4. It is easy to see that the orientation τ of S which is induced by the entire vertex set of the grid is a tangle. Let us show that τ is $\Omega(n^2)$-resilient. Every element (A, B) of τ satisfies $|A| \leqslant 10$; indeed, most satisfy $|A| \leqslant 5$. Thus, any set of separations in τ with a co-small supremum has at least $n^2/10$ elements as the $(n \times n)$-grid has precisely n^2 vertices.

Why can the notion of resilience help us with constructing a function that induces a given orientation? Consider an orientation τ of a set S of separations of a set V, and write $\mu = \mu(\tau)$ for the set of maximal elements of τ. Let us see how resilience that is large in terms of $|\mu|$ can help us build a τ-inducing function.

Assume that τ is k-resilient for some $k \in \mathbb{N}$, and recall that $\mu^{(k)}$ denotes the set of k-element subsets of μ. Then the resilience of τ implies that for every $\mu' \in \mu^{(k)}$, there exists some $v_{\mu'} \in V$ that is not contained in the small side of any separation in μ'; in particular, the set $\{v_{\mu'}\}$ induces μ'. It seems natural to construct a function that induces μ (and thus τ, by Lemma 5.2.1) by combining all these local μ'-inducing sets $\{v_{\mu'}\}$, i.e. by assigning to each $v \in V$ as its weight the number of sets $\mu' \in \mu^{(k)}$ with $v_{\mu'} = v$.

It turns out that the weight function w on V defined in this way need not in general be a function that induces μ. This is because each $v_{\mu'}$, while adding its weight to the big sides of the separations in μ', can also add weight to the small sides of separations in $\mu \smallsetminus \mu'$. But as soon as k is large enough in that each fixed separation $(A, B) \in \mu$ is contained in the majority of the sets in $\mu^{(k)}$, which will happen as soon as μ has more $(k-1)$-subsets to form a k-subset with (A, B) than it has k-subsets not including (A, B), the orientation (A, B)

of $\{A, B\}$ will be induced by the majority of the sets $\{v_{\mu'}\}$ that locally induce $\mu' \in \mu^{(k)}$. We can then deduce from this that w induces μ and hence τ by Lemma 5.2.1.

More precisely, we have the following generalisation of Theorem 5.1.1 to arbitrary orientations τ:

Theorem 5.3.1. *Let S be a set of separations of a set V, and let τ be an orientation of S. Let m be the number of maximal elements of τ. If τ is k-resilient for some integer $k > \frac{m}{2}$, then τ is induced by a function on V.*

We will formally obtain Theorem 5.3.1, and hence Theorem 5.1.1, below as a corollary of the more general Theorem 5.3.5. But before we do so, let us first show that both Theorem 5.1.1 and Theorem 5.3.1 are optimal with respect to the parameter k in k-resilience. To show this we exhibit a suitable tangle that is not induced by a function, by using a more general version of the construction from the introduction.

Proposition 5.3.2. *For all $m, k \in \mathbb{N}$ with $3 \leqslant k \leqslant \frac{m}{2}$, there exist a set V, a submodular order function $|\cdot|$ on $U_{bip}(V)$, and an m-tangle $\tau_{m,k}$ in $U_{bip}(V)$ that has m maximal elements and is k-resilient, but which is not induced by any function on V.*

An example of the tangles in Proposition 5.3.2 is given by a certain type of hypergraph edge tangles introduced in [EKT20]. We now describe their construction, and then show that these tangles do indeed have all the desired properties.

Proof of Proposition 5.3.2. Let $V = [m]^{(k)}$ consist of all k-element subsets of $[m]$. For every $i \in [m]$ let $V_i = \{X \in V \mid i \in X\}$ be the set of all k-element subsets of $[m]$ containing i. We assign to each bipartition $\{A, B\}$ of V as its order $|A, B|$ the number of sets V_i meeting both A and B. This order function $|\cdot|$ on $U_{bip}(V)$ is easily seen to be submodular (see [EKT20] for a formal proof).

Every bipartition $\{A, B\}$ of V of order less than m has precisely one side which contains V_i for some $i \in [m]$. Indeed, one such side exists by the definition of the order function $|\cdot|$, and since $V_i \cap V_j \neq \emptyset$ for distinct $i, j \in [m]$, this side is unique. Hence,

$$\tau_{m,k} := \{(A, B) \mid \{A, B\} \in S_m \text{ and } \exists\, i \in [m] \colon V_i \subseteq B\}$$

is a well-defined orientation of $S_m \subseteq U_{bip}(V)$. The maximal elements of $\tau_{m,k}$ are precisely the $\vec{s_i} = (V \smallsetminus V_i, V_i)$ for $i \in [m]$. So in order to see that $\tau_{m,k}$ is indeed an m-tangle in $U_{bip}(V)$, it is enough to observe that $\vec{s_i} \vee \vec{s_j} \vee \vec{s_\ell}$ is not co-small for any $1 \leqslant i, j, \ell \leqslant m$. But since $k \geqslant 3$, we always have $V_i \cap V_j \cap V_\ell \neq \emptyset$.

As shown in [EKT20], it is immediate from double counting (as in the example from the introduction) that $\tau_{m,k}$ is not induced by any function on V for $m \geqslant 6$ and $k \leqslant \frac{m}{2}$. So it remains to check that $\tau_{m,k}$ is k-resilient. But this is immediate from the construction: for every collection $\vec{s_{i_1}}, \dots, \vec{s_{i_k}}$ of k distinct maximal elements of $\tau_{m,k}$, we have $\{i_1, \dots, i_k\} \in \bigcap_{j \in [k]} V_{i_j}$. Hence for every collection of at most k maximal elements of $\tau_{m,k}$, the intersection of their big sides is non-empty. □

The $\tau_{m,k}$ constructed in our proof of Proposition 5.3.2 are abstract tangles, but the construction can easily be modified to find $\mathcal{F}^\ell$-tangles for arbitrary $\ell > 1$ with the same properties: instead of taking the k-element subsets of $[m]$ as the set V, we can take V as the disjoint union of $\lceil \ell \rceil$-element sets, one for every k-element subset of $[m]$.

Before we proceed towards a proof of Theorem 5.3.5, our generalisation of Theorem 5.3.1, let us briefly investigate in some more detail the above examples of $\mathcal{F}^\ell$-tangles that are not induced by functions. Note that our construction of $\mathcal{F}^\ell$-tangles does not necessarily work when the value of ℓ is not constant, but large in terms of $|V|$, e.g. of size at least $\varepsilon|V|$ for some constant $\varepsilon > 0$. The following proposition shows that there exists a sharp lower bound for those $\varepsilon > 0$ for which $\ell \geqslant \varepsilon|V|$ guarantees the existence of a function that induces the $\mathcal{F}^\ell$-tangle.

Proposition 5.3.3. *Let V be an n-set, and let $0 < \varepsilon < 1$. If $\varepsilon \geqslant 1/8$, then every $\mathcal{F}^{\varepsilon n}$-tangle τ of a set S of bipartitions of V is induced by a function on V; if $\varepsilon > 1/8$, then τ is even induced by a subset of V.*

Conversely, for every $\varepsilon < 1/8$ there exist $n \in \mathbb{N}$ and a set S of bipartitions of an n-set such that some $\mathcal{F}^{\varepsilon n}$-tangle of S is not induced by any function on V.

Proof. Let τ be an $\mathcal{F}^\ell$-tangle of a set S of bipartitions of a set V satisfying $\ell \geqslant |V|/8$. If all of V induces τ, then we are done; so suppose not. Then there exists $(A_1, B_1) \in \tau$ with $|B_1| \leqslant |V|/2$. Again we are done if B_1 induces τ. If this is not the case, then there is some $(A_2, B_2) \in \tau$ with $|B_1 \cap A_2| \geqslant |B_1 \cap B_2|$; in particular, we have $|B_1 \cap B_2| \leqslant |V|/4$.

It turns out that if $\ell > |V|/8$, then $B_1 \cap B_2$ has to induce τ, since otherwise there exists some $(A_3, B_3) \in \tau$ such that

$$|(B_1 \cap B_2) \cap B_3| \leqslant |(B_1 \cap B_2) \cap A_3|.$$

This implies $|(B_1 \cap B_2) \cap B_3| \leqslant |V|/8$, which contradicts the fact that τ is an $\mathcal{F}^\ell$-tangle.

If $\ell = |V|/8$, then the same arguments as above produce a τ-inducing set if at least one of the occurring inequalities is strict. So suppose that all the above inequalities are satisfied with equality. In particular, every $(A, B) \in \tau$ satisfies $|A| \leqslant |B|$.

Using similar reasoning as above, we can obtain a function that induces τ. It will be enough to find a weight function w on V that induces the set $\tau' \subseteq \tau$ consisting of all $(A, B) \in \tau$ with $|A| = |B|$. For a given function w inducing τ', we can obtain a function inducing τ by adding large enough constant weight to all the vertices in V.

Suppose there are $(A, B), (C, D) \in \tau'$ with $|B \cap C| > |B \cap D|$. This yields $|B \cap D| < |V|/4$, which in turn implies the existence of a set inducing τ, by the same arguments as above. Consequently, for every $(A, B), (C, D) \in \tau'$, we have $|B \cap C| \leqslant |B \cap D|$.

Hence, the weight function w defined by counting for every $v \in V$ the number of separations $(A, B) \in \tau'$ with $v \in B$ is a function inducing τ': given $(C, D) \in \tau'$, we have $w(C) = \sum_{(A,B)\in\tau'} |B \cap C|$ and $w(D) = \sum_{(A,B)\in\tau'} |B \cap D|$. As above, we obtain that $|B \cap C| \leqslant |B \cap D|$ for every $(A, B) \in \tau'$, and as τ is an $\mathcal{F}^\ell$-tangle, we clearly have $D \neq \emptyset$ and hence $|D \cap C| < |D \cap D|$. This yields $w(C) < w(D)$, and thus, w induces τ.

For the second part of the proposition, consider V and the m-tangle $\tau_{m,k}$ in $U_{bip}(V)$ as constructed in our proof of Proposition 5.3.2 for some $m \geqslant 2k \geqslant 6$. Then for any three maximal elements of $\tau_{m,k}$, their intersection contains exactly $\binom{m-3}{k-3}$ elements of V. In particular, $\tau_{m,k}$ is an $\mathcal{F}^\ell$-tangle for all $\ell < \binom{m-3}{k-3}$.

Recall that $|V| = \binom{m}{k}$. For $k = m/2$ we have $\lim_{m\to\infty} \binom{m-3}{k-3}/\binom{m}{k} = 1/8$. Thus, we find for any $\varepsilon < 1/8$ some $n = \binom{m}{k} \in \mathbb{N}$ such that the tangle $\tau_{m,k}$ from Proposition 5.3.2 is an $\mathcal{F}^{\varepsilon n}$-tangle of a set of bipartitions of an n-set that is not induced by any function on V. □

Back to Theorem 5.3.1, recall that this theorem is sharp in terms of the parameter k in k-resilience as shown by Proposition 5.3.2. However, the converse of Theorem 5.3.1 fails, i.e. not even every tangle induced by a set of points has high resilience compared to the number of its maximal elements.

Example 5.3.4. Let V be a set of size $n \geqslant 4$, and let S be the set of all bipartitions of V that have a side of size less than $n/3$. Let τ be the orientation of S which orients every $s \in S$ in such a way that its big side contains more elements of V than its small side. In particular, V induces τ.

This orientation τ of S is a tangle, since no three big sides of bipartitions in τ have empty intersection. However, four big sides can, so the supremum of four bipartitions in τ can be co-small. Thus, τ has resilience 3.

Now τ has $m = \binom{n}{\lceil (n/3)-1 \rceil}$ maximal elements, namely those bipartitions of V whose small side has maximum size. In particular, the resilience of τ is low compared with m, although τ is induced by a function on V, and even by the set V.

It turns out that we can generalise the notion of resilience in a way which includes the tangle from the previous Example 5.3.4 without invalidating Theorem 5.3.1. In fact, our more general notion leads to a more general result, Theorem 5.3.5, which actually characterises the orientations induced by functions.

Our more general notion of resilience is based on our earlier observation that, if τ is k-resilient, then this provides for every $\mu' \in \mu^{(k)}$ a singleton set $\{v_{\mu'}\}$ that induces μ'. In the following definition we ask, instead of k-resilience, that for every $\mu' \in \mu^{(k)}$ there exists a function $w_{\mu'}$ which induces μ' and is not too badly wrong on the separations in $\mu \setminus \mu'$.

More precisely, an orientation τ of a set S of separations of a set V is *k-locally ℓ-induced* for given $k \in \mathbb{N}$ and $\ell \geqslant 0$ if for every set $\mu' \subseteq \tau$ of size $|\mu'| \leqslant k$, there is a weight function $w_{\mu'}$ on V satisfying

(i) $\forall\ (A, B) \in \mu' \colon w_{\mu'}(B) - w_{\mu'}(A) \geqslant 1$;

(ii) $\forall\ (A, B) \in \tau \colon w_{\mu'}(A) - w_{\mu'}(B) \leqslant \ell$.

Observe that by Lemma 5.2.1 one needs only consider sets $\mu' \subseteq \mu$ in the above definition where $\mu = \mu(\tau)$ is the set of maximal elements of τ. In addition, we can equivalently strengthen the condition of $|\mu'| \leqslant k$ to $|\mu'| = k$ (i.e. $\mu' \in \mu^{(k)}$), as long as τ has at least k maximal elements.

The above definition is indeed a generalisation of our earlier notion of resilience, since a k-resilient orientation τ of S is k-locally 1-induced: for $\mu' \subseteq \tau$ with $|\mu'| \leqslant k$, take $w_{\mu'}$ assigning 1 to a single element in $V \setminus \bigcup_{(A,B)\in\mu'} A$, which is non-empty since τ is k-resilient, and 0 to all other elements in V.

If τ is induced by a function on V, then τ is k-locally ℓ-induced for all $k \in \mathbb{N}$ and $\ell \geqslant 0$. Indeed, as described in Section 5.2.3, there exists a function that induces τ and satisfies (i) for all separations in τ, and any function inducing τ clearly satisfies (ii). In particular, the tangle in Example 5.3.4 is k-locally ℓ-induced for all $k \in \mathbb{N}$ and $\ell \geqslant 0$.

Here, then, is our generalisation of Theorem 5.3.1:

Theorem 5.3.5. *Let τ be an orientation of a set S of separations of a set V, and suppose that τ has m maximal elements. Then τ is induced by a function on V if and only if it is k-locally ℓ-induced for some $k \in \mathbb{N}$ and $\ell > 0$ with $k > \frac{m}{1+1/\ell}$.*

Since every k-resilient orientation τ is k-locally 1-induced, Theorem 5.3.1, and hence Theorem 5.1.1 as well, are direct corollaries of Theorem 5.3.5.

Proof of Theorem 5.3.5. We have seen above that if τ is induced by a function on V, then it is k-locally ℓ-induced for every $k \in \mathbb{N}$ and $\ell \geqslant 0$. In particular, τ is m-locally ℓ-induced for every $\ell > 0$, and we have $m > \frac{m}{1+1/\ell}$ in this case.

For the converse, recall that by Lemma 5.2.1 it is enough to show that the set $\mu = \mu(\tau)$ of maximal elements of τ is induced by some function on V. If $k \geqslant m$ the statement is true immediately; so suppose for the following that $k < m$. We construct a function w inducing μ as follows: for every $\mu' \in \mu^{(k)}$ let $w_{\mu'}$ be a weight function on V as in the definition of 'k-locally ℓ-induced'. Then we combine all these $w_{\mu'}$ to define the weight function w on V as

$$w : V \to \mathbb{R}_{\geqslant 0}, w(v) = \sum_{\mu' \in \mu^{(k)}} w_{\mu'}(v).$$

We show that w induces μ, as desired.

For an arbitrary separation $\vec{s} = (A, B) \in \mu$, let $\mu^{(k)}_{\vec{s}}$ consist of all $\mu' \in \mu^{(k)}$ containing $\vec{s}$. Then $\mu^{(k)}_{\vec{s}}$ has size $\binom{m-1}{k-1}$ since $\vec{s}$ is contained in $\binom{m-1}{k-1}$ many sets $\mu' \in \mu^{(k)}$. Similarly, $\mu^{(k)} \setminus \mu^{(k)}_{\vec{s}}$ has size $\binom{m-1}{k}$. Therefore, (i) in the definition of 'k-locally ℓ-induced' yields

$$\sum_{\mu' \in \mu^{(k)}_{\vec{s}}} w_{\mu'}(A) \leqslant \sum_{\mu' \in \mu^{(k)}_{\vec{s}}} (w_{\mu'}(B) - 1) = \sum_{\mu' \in \mu^{(k)}_{\vec{s}}} w_{\mu'}(B) - \binom{m-1}{k-1}.$$

Similarly, we obtain by (ii) that

$$\sum_{\mu' \in \mu^{(k)} \setminus \mu^{(k)}_{\vec{s}}} w_{\mu'}(A) \leqslant \sum_{\mu' \in \mu^{(k)} \setminus \mu^{(k)}_{\vec{s}}} (w_{\mu'}(B) \ + \ \ell) = \sum_{\mu' \in \mu^{(k)} \setminus \mu^{(k)}_{\vec{s}}} w_{\mu'}(B) \ + \ \ell \ \cdot \ \binom{m-1}{k}.$$

These inequalities combine to

$$\begin{aligned} w(A) &= \sum_{\mu' \in \mu^{(k)}} w_{\mu'}(A) \\ &= \sum_{\mu' \in \mu^{(k)}_{\vec{s}}} w_{\mu'}(A) + \sum_{\mu' \in \mu^{(k)} \setminus \mu^{(k)}_{\vec{s}}} w_{\mu'}(A) \\ &\leqslant \sum_{\mu' \in \mu^{(k)}_{\vec{s}}} w_{\mu'}(B) - \binom{m-1}{k-1} + \sum_{\mu' \in \mu^{(k)} \setminus \mu^{(k)}_{\vec{s}}} w_{\mu'}(B) + \ell \cdot \binom{m-1}{k} \\ &= w(B) - \binom{m-1}{k-1} + \ell \cdot \binom{m-1}{k}. \end{aligned}$$

Now since $k > \frac{m}{1+1/\ell}$ and $k < m$, we have that $\ell < \frac{k}{m-k}$ and thus

$$\ell \cdot \binom{m-1}{k} < \binom{m-1}{k-1},$$

as $\binom{m-1}{k} = \frac{(m-1)!}{k!(m-1-k)!}$ and $\binom{m-1}{k-1} = \frac{(m-1)!}{(k-1)!(m-k)!}$ differ precisely by the factor $\frac{k}{m-k}$. This then implies $w(A) < w(B)$, so w induces $\vec{s} = (A, B)$. Thus, since $\vec{s} \in \mu$ was arbitrarily chosen, w induces μ and hence τ, by Lemma 5.2.1. □

5.4 Orientations and duality of set separations

In this section we present a second characterisation of those orientations of a set of separations that are induced by a function on the ground set: a characterisation in terms of a duality between separation systems. As an unexpected corollary, we obtain an independent second proof of Theorem 5.3.1.

The duality of separation systems, which was introduced in [D19] and first studied in [D21; DEET23], is defined for set separations as follows. Let S be a set of separations of a set V, and let σ be an orientation of S. For every $v \in V$, the sets

$$C_\sigma(v) := \{\{A, B\} \in S \mid (A, B) \in \sigma, v \in A\} \text{ and}$$
$$D_\sigma(v) := \{\{A, B\} \in S \mid (A, B) \in \sigma, v \in B\}$$

form the sides of the separation $\{C_\sigma(v), D_\sigma(v)\}$ of S. The map $\varphi_\sigma : V \to U(S)$ then associates with $v \in V$ the separation $\varphi_\sigma(v) := \{C_\sigma(v), D_\sigma(v)\}$ of S. We write $V_\sigma := \varphi_\sigma(V)$ for the *dual set of separations of S with respect to σ*.

Suppose now that φ_σ is injective.[4] Then the dual set V_σ of separations of S with respect to σ has a natural *default orientation* τ_σ which orients every $\varphi_\sigma(v) = \{C_\sigma(v), D_\sigma(v)\} \in V_\sigma$ as $(C_\sigma(v), D_\sigma(v))$, i.e. its big side contains those $\{A, B\} \in S$ with $(A, B) \in \sigma$ and $v \in B$. It is easy to see that the dual set of separations of V_σ with respect to τ_σ is again S with default orientation σ. In simple terms, 'dualising the dual yields the primal' [DEET23].

If φ_σ is injective, we can ask whether the existence of a function on V that induces the orientation σ of S relates to any property of the default orientation τ_σ of the dual set V_σ of separations of S with respect to σ. It turns out that it does, and it does so in an intriguing way: σ is induced by a function on V if and only if every non-zero weight function w' on S induces some separation in τ_σ. More generally, we have the following theorem:

Theorem 5.4.1. *Let S be a set of separations of a set V, and let σ be an orientation of S. Then the following two assertions are equivalent:*

(i) *There exists a function on V that induces σ.*

(ii) *For every non-zero weight function w' on S, there exists $v \in V$ such that $w'(C_\sigma(v)) < w'(D_\sigma(v))$.*

In particular, if φ_σ is injective and τ_σ is the default orientation of the dual set V_σ of separations of S with respect to σ, then σ is induced by a function on V if and only if every non-zero weight function w' on S induces a separation in τ_σ.

[4]The map φ_σ is injective if and only if for every two distinct elements $v, v' \in V$, there exist two separations in S for one of which v and v' are contained in the same side and for the other in different sides. If this is not the case for $v, v' \in V$, then v and v' 'carry the same information' about how the separations in S separate V and can hence be seen as redundant.

The proof of Theorem 5.4.1 will be done purely in terms of linear algebra and can be followed without any further knowledge about our particular duality of separation systems described above. The key tool in our proof will be the following variant of Farkas's Lemma (see, e.g., [P17, 6. Theorem][5]).

Lemma 5.4.2 (Farkas's Lemma)**.** *Let $Q \in \mathbb{R}^{n\times\ell}$ and $b \in \mathbb{R}^\ell$. Then exactly one of the following two assertions holds:*

(i) *There exists $x \in \mathbb{R}^n_{\geqslant 0}$ with $Q^T x \geqslant b$.*

(ii) *There exists $y \in \mathbb{R}^\ell_{\geqslant 0}$ with $Qy \leqslant 0$ and $b^T y > 0$.*

Proof of Theorem 5.4.1. Fix enumerations $V = \{v_1, \dots v_n\}$ and $S = \{s_1, \dots, s_\ell\}$, and let $\vec{s_j} = (A_j, B_j) \in \sigma$ for $j \in [\ell]$. Using these enumerations, we shall, for the course of this proof, identify a weight function w on V with a vector $x = x(w)$ in $\mathbb{R}^n_{\geqslant 0}$ and a weight function w' on S with a vector $y = y(w')$ in $\mathbb{R}^\ell_{\geqslant 0}$.

Let us define a matrix $Q = Q(\sigma) \in \mathbb{R}^{n\times\ell}$ via

$$Q_{ij} = \begin{cases} 1, & v_i \in B_j \smallsetminus A_j; \\ 0, & v_i \in A_j \cap B_j; \\ -1, & v_i \in A_j \smallsetminus B_j. \end{cases}$$

Recall from the definition of φ_σ that $C_i := C_\sigma(v_i)$ (respectively $D_i := D_\sigma(v_i)$) consists of all those $\{A_j, B_j\} \in S$ with $v_i \in A_j$ (respectively $v_i \in B_j$). So given a weight function w' on S, we obtain for $y = y(w') \in \mathbb{R}^\ell_{\geqslant 0}$ and all $i \in [n]$ that

$$(Qy)_i = \sum_{B_j \ni v_i} y_j - \sum_{A_j \ni v_i} y_j = w'(D_i) - w'(C_i).$$

So a non-zero weight function w' on S is *not* as in Theorem 5.4.1 (ii) if and only if $Qy \leqslant 0$ for $y = y(w') \in \mathbb{R}^\ell_{\geqslant 0}$ where $\leqslant$ is meant coordinate-wise.

Similarly, let w be a weight function on V and $x = x(w) \in \mathbb{R}^n_{\geqslant 0}$. Then we compute for all $j \in [\ell]$ that

$$(Q^T x)_j = \sum_{v_i \in B_j} x_i - \sum_{v_i \in A_j} x_i = w(B_j) - w(A_j).$$

So a weight function w on V induces σ if and only if we have $Q^T x > 0$ for $x = x(w) \in \mathbb{R}^n_{\geqslant 0}$. Recall from Section 5.2.3 that σ is induced by some function on V if and only if there exists a function w with $w(B) - w(A) \geqslant 1$ for all $(A, B) \in \sigma$. Hence, σ is induced by some function on V (as in Theorem 5.4.1 (i)) if and only if $Q^T x \geqslant \mathbf{1}$ where $x = x(w) \in \mathbb{R}^n_{\geqslant 0}$ for some non-zero weight function w on V and $\mathbf{1}$ is the constant 1 vector in $\mathbb{R}^\ell$.

The result then follows by applying Lemma 5.4.2 to Q and $b = \mathbf{1} \in \mathbb{R}^\ell$, and denoting $x = x(w)$ in Lemma 5.4.2 (i) and $y = y(w')$ in Lemma 5.4.2 (ii). The 'in particular'-part is immediate from the definition of V_σ and τ_σ. □

[5] Our version of Farkas's Lemma follows from [P17, 6. Theorem] by applying their theorem to $A = Q^T$ and $-b$ instead of b.

As mapping a separation in S to its orientation in σ is a bijection between S and σ, we could equivalently define the weight function w' in Theorem 5.4.1 (ii) on the orientation σ of S; for notational simplicity we will freely switch between these two definitions in what follows.

By Lemma 5.2.1, an orientation σ of S is induced by a function on V if and only if its set of maximal elements is. So applying Theorem 5.4.1 to the set $\mu = \mu(\sigma)$ of maximal elements of σ and its underlying set of unoriented separations, we have the following corollary.

Corollary 5.4.3. *Let S be a set of separations of a set V. Let σ be an orientation of S, and let $\mu = \mu(\sigma)$ be the set of maximal elements of σ. Then the following two assertions are equivalent:*

(i) *There exists a function on V that induces σ.*

(ii) *For every non-zero weight function w' on μ, there exists $v \in V$ with*

$$\sum_{\substack{(A,B)\in\mu \\ \text{with } v\in A}} w'((A,B)) < \sum_{\substack{(A,B)\in\mu \\ \text{with } v\in B}} w'((A,B)). \qquad \square$$

As an illustration of the power of Corollary 5.4.3 and hence Theorem 5.4.1, let us re-prove Theorem 5.3.1, which asserts that highly resilient orientations are induced by functions on the ground set.

Proposition 5.4.4. *Let S be a set of separations of a set V, and let σ be an orientation of S with m maximal elements. If σ is k-resilient for some integer $k > \frac{m}{2}$, then σ is induced by a function on V.*

Proof. Let $\mu = \mu(\sigma)$ be the set of maximal elements of σ. We apply Corollary 5.4.3 to σ in that we consider an arbitrary non-zero weight function w' on μ and show that case (ii) in Corollary 5.4.3 holds. Let $\mu' \subseteq \mu$ consist of those k separations in μ which have the highest weight with respect to w'. Since we have $k > \frac{m}{2}$ by assumption, this immediately yields $w'(\mu') > w'(\mu \setminus \mu')$.

Now σ is k-resilient, so there exists some $v \in V$ which is not contained in the small side of any separation in μ'. By the choice of μ', this v satisfies

$$\sum_{\substack{(A,B)\in\mu \\ \text{with } v\in B}} w'((A,B)) \geqslant w'(\mu') > w'(\mu \setminus \mu') \geqslant \sum_{\substack{(A,B)\in\mu \\ \text{with } v\in A}} w'((A,B)).$$

Thus, Corollary 5.4.3 (ii) holds for w', and since w' was arbitrarily chosen, this implies that σ is induced by a function on V. $\square$

5.5 Extendable tangles are induced by point sets

Let S be a set of separations of a set V of 'points', and let τ be an orientation of S. In Section 5.3 we analysed various properties of τ which ensure that τ is induced by a function on V. All these properties required us to consider large subsets of τ instead of the usual

triples which are required for the definition of a tangle of S. In particular, all the notions considered above may be viewed as strengthenings of the triple condition in the definition of a tangle, i.e., we give a stronger condition that an orientation needs to satisfy in order to be a tangle that is induced by some function on V.

But how can we guarantee the existence of a function inducing a tangle τ of S if we do not want to strengthen the definition of a tangle in the above sense? We know that there exist tangles that are not induced by any function on the ground set of the separations they orient (see e.g. Proposition 5.3.2). So instead of looking for a function inducing τ itself, we may try to find one that induces some $\tau' \subseteq \tau$. Ideally, we can do this in such a way that the function inducing τ' is still, in some sense, related to the original tangle τ.

Given an order function on a universe U of separations of a set V, one natural such subset of a k-tangle τ in U, say, consists of all separations in τ of order less than some $k' < k$. In other words, we would like to obtain, given a k-tangle τ in U, a function on V that induces the k'-tangle $\tau' \subseteq \tau$ in U.

One way in which we could try to achieve this consists in proving the following: there exists a function $f : \mathbb{N} \to \mathbb{N}$ with $f(k) \geqslant k$ for all $k \in \mathbb{N}$ such that if a k-tangle τ' in U extends to a $f(k)$-tangle τ in U, then τ' is induced by a function on V, or even by a subset of V. In this case, we may view the function w inducing τ' as an approximation of a function inducing its extension τ – although w will in general not orient $S_{f(k)} \subseteq U$ as τ.

Consider for example the m-tangle $\tau_{m,k}$ in $U_{bip}(V)$ constructed in Proposition 5.3.2 for some $3 \leqslant k \leqslant \frac{m}{2}$. This tangle $\tau_{m,k}$ is not induced by any function on V, but if we consider only the separations of order less than $\frac{m}{2}$ in this example, then they are even induced by a set: V orients all the separations in $U_{bip}(V)$ of order less than $\frac{m}{2}$ in the same way as $\tau_{m,k}$.

This leads us to the question of whether tangles which extend to tangles of twice their order are always induced by functions or even sets, i.e., whether $f(k) := 2k$ is suitable. We show that this is indeed the case in that k-profiles in U which extend to regular $2k$-profiles in U are induced by subsets of V – as long as we work in the universe $U = U(V)$ of separations of a set V equipped with our standard order function on U which assigns to a separation the cardinality of its separator as its order. Recall that, for this order function, all k-profiles in U are induced by functions on V [EKT20], but those as above are even induced by subsets of V:

Theorem 5.5.1. *Let $U = U(V)$ be the universe of all separations of a set V and let $|\cdot|$ be the standard order function on U. If τ' is a k-profile in U for some $k \in \mathbb{N}$ that extends to a regular $2k$-profile τ in U, then τ' is induced by a set $X \subseteq V$ of size at least $2k$.*

Proof. If V has less than $2k$ elements, then there exists no regular $2k$-profile τ in U, since $(V, V) \in \tau$ contradicts its regularity. Thus, the theorem always holds for $|V| < 2k$, and we may assume $|V| \geqslant 2k$ in what follows.

Our desired set X inducing τ' will be the *interior* $\bigcap_{(A,B)\in\sigma} B$ of a star σ contained in τ. Let us first show that such interiors cannot be too small.

Claim 1. *The interior of any star contained in a regular $2k$-profile in U has at least $2k$ elements.*

Proof. Suppose not, let τ be a regular $2k$-profile in U, and let $\sigma \subseteq \tau$ be a star whose

interior $X = \bigcap_{(A,B)\in\sigma} B$ has size $|X| < 2k$. Note that σ is non-empty as the interior of the empty star is the whole set V which by assumption has size at least $2k$.

Let us write $\sigma = \{(A_1, B_1), \dots, (A_\ell, B_\ell)\}$. We claim that $|(A_1, B_1) \vee \cdots \vee (A_i, B_i)| < 2k$ for any $i \leqslant \ell$. By definition, we have

$$|(A_1, B_1) \vee \cdots \vee (A_i, B_i)| = |(A_1 \cup \cdots \cup A_i) \cap (B_1 \cap \cdots \cap B_i)|.$$

Since σ is a star, we have $(A_1 \cup \cdots \cup A_i) \subseteq B_j$ for every $j > i$. So in particular, we have

$$(A_1 \cup \cdots \cup A_i) \cap (B_1 \cap \cdots \cap B_i) \subseteq B_1 \cap \cdots \cap B_\ell = X.$$

Therefore, $|(A_1, B_1) \vee \cdots \vee (A_i, B_i)| \leqslant |X| < 2k$. Since τ is a profile, it follows inductively that $(A_1, B_1) \vee \cdots \vee (A_i, B_i) \in \tau$ for every $i \leqslant \ell$. In particular, $(Y, X) = (A_1, B_1) \vee \cdots \vee (A_\ell, B_\ell) \in \tau$ where $Y = \bigcup_{(A,B)\in\sigma} A$.

Since $|X| < 2k$, the separation $\{X, V\}$ has order $< 2k$ and hence an orientation in τ. By the regularity of τ, this orientation must be (X, V) because (V, X) is co-small. But this leads to a contradiction since this would imply $(Y, X) \vee (X, V) = (V, X) \in \tau$ as τ is a profile. □

Let $\sigma \subseteq \tau$ be a star whose interior $X = \bigcap_{(C,D)\in\sigma} D$ is of smallest size among all stars contained in τ. By Claim 1 we have $|X| \geqslant 2k$. We claim that X induces τ'.

To prove this, we show that $|X \cap A| < k$ for every $(A, B) \in \tau'$. Since we have $|X| \geqslant 2k$, this immediately implies that X induces τ'. So suppose for a contradiction that there exists $(A, B) \in \tau'$ with $|X \cap A| \geqslant k$, and assume that (A, B) has minimal order among all such separations in τ'. Note that (A, B) may be induced by X.

Now for every $(C, D) \in \sigma$, the separation $(A \cap D, B \cup C)$ has at least the order of (A, B). Indeed, we have $X \subseteq D$ by construction, and therefore we get $|(A \cap D) \cap X| = |A \cap X| \geqslant k$. Moreover, if $(A \cap D, B \cup C)$ would have order less than (A, B), then $(A \cap D, B \cup C)$ would be contained in τ (since τ is a profile) and hence in $\tau' \subseteq \tau$ because (A, B) has order less than k. Then, however, $(A \cap D, B \cup C)$ would contradict the minimal choice of (A, B).

By the submodularity of the standard order function, the order of $(B \cap C, A \cup D)$ is at most $|C, D|$. Since $(B \cap C, A \cup D) \leqslant (C, D) \in \tau$ and τ is consistent, we thus have $(B \cap C, A \cup D) \in \tau$. Hence, the star

$$\hat{\sigma} := \{(A, B)\} \cup \{(B \cap C, A \cup D) \mid (C, D) \in \sigma\}$$

is contained in τ.

We claim that the interior $\hat{X}$ of $\hat{\sigma}$ is smaller than X contradicting the choice of σ. Indeed, by definition, we have

$$\hat{X} = B \cap \bigcap_{(C,D)\in\sigma} (A \cup D) = (A \cap B) \cup (B \cap X) = ((A \cap B) \smallsetminus X) \cup (B \cap X).$$

Since (A, B) is a separation of V, the set $X \subseteq V$ is the disjoint union of $B \cap X$ and $(A \cap X) \smallsetminus B$. So we are done if

$$|(A \cap B) \smallsetminus X| < |(A \cap X) \smallsetminus B|.$$

Let $h = |A \cap B \cap X|$. Since $|A \cap X| \geqslant k$, we have $|(A \cap X) \smallsetminus B| \geqslant k - h$. However, we have $(A, B) \in \tau$, so $|A \cap B| < k$ and hence $|(A \cap B) \smallsetminus X| < k - h$, completing the proof. □

Our proof of Theorem 5.5.1 heavily relies on the assumption that the order of a separation in $U(V)$ is given by the size of its separator. We do not know if a similar result holds for other or even all submodular order functions on $U(V)$.

Problem 5.5.2. *Let V be a set, and consider any submodular order function on $U(V)$. Is it true that if τ' is a k-profile in $U(V)$ for some $k \in \mathbb{N}$ which extends to a regular $2k$-profile τ in $U(V)$, then τ' is induced by some subset of V? What happens for other universes of separations of V such as $U_{bip}(V)$?*

CHAPTER 6

Efficiently distinguishing all tangles in locally finite graphs

This chapter is based on [JK24] and joint work with Paul Knappe.

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.

6.1 Introduction

In their Graph Minors Project [RS83–04], 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 6.1.1 ([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 [D18b; 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 6.1.1 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 [EKT22b, Example 4.9][1]. On the other hand, they and, independently, Carmesin, Hamann and Miraftab 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 6.1.2 ([CHM22, Theorem 7.3] & [EKT22b, Theorem 6.6])**.** *Every connected locally finite graph has a canonical tree of tangles.*

So in order to understand to which locally finite graphs Theorem 6.1.1 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 6.1.2, 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 16. *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 16 does not hold, as the infinite grid witnesses. More generally, [DJKK26, Lemma 8.5] implies that every connected locally finite graph that is quasi-transitive and accessible admits a tree-decomposition as in Theorem 16 – no matter if it has a thick end or not.

For a proof of Theorem 16, it suffices to show that some (canonical) tree of tangles, which exists by Theorem 6.1.2, 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.

[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.

Elbracht, Kneip and Teegen [EKT22b, 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 16 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 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 6.4.1, 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 17. *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 6.2. In Section 6.3 we study limit separations and develop various general techniques for their investigation. Then we introduce our interlacing theorem in Section 6.4. Finally, we combine our results on limit separations with the interlacing theorem to prove our main results, Theorems 16 and 17, in Section 6.5.

6.2 Preliminaries

For the general definitions and notations around graphs, we follow [D17b]. 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:

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

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 6.2.1 ([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.

6.3 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.

6.3.1 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 6.3.1. *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 6.3.2. *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 $X \subseteq A \smallsetminus B = V(G) \smallsetminus B$, there exists $I \in \mathbb{N}$ such that $X \subseteq V(G) \smallsetminus B_i = A_i \smallsetminus B_i$ 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 6.3.3. *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 $A \smallsetminus B$ equal the respective intersections with $A_i \cap B_i$ and $A_i \smallsetminus B_i$.* □

6.3.2 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 [EKT22b, 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 16 to obtain a tree-decomposition as in Theorem 16 from the nested set given by Theorem 6.1.2.

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 $B = \emptyset$ [DJKK26, 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 6.3.4 ([EKT22b, 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 6.3.2 provides us a tool for this: when we investigate the limit of a strictly increasing sequence in $\overrightarrow{N}$, then Lemma 6.3.2 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 6.3.2 in combination with the next Lemma 6.3.5 in the paragraph after stating the latter.

To state Lemma 6.3.5, 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. [EKT22b, 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 6.3.5 ([DJKK26, 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 6.3.2 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 6.3.2 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 6.3.5 which contradicts by Lemma 6.3.2 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 6.3.6. *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 6.3.6, we show two auxiliary results. The first lemma observes that $A \cap B$ is non-empty even under weaker assumptions.

Lemma 6.3.7. *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 6.3.8. *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 6.3.2 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 6.3.8 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 6.3.7 and Lemma 6.3.8 at hand, we are ready to prove Theorem 6.3.6.

Proof of Theorem 6.3.6. Suppose for a contradiction that $A \cap B$ is finite. Then Lemma 6.3.3 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 6.3.2 and Lemma 6.3.5 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 6.3.7, and Lemma 6.3.8 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 6.3.6, 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 6.5, this is already the first step towards Theorem 17 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 6.5.

6.3.3 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 6.5, 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 6.3.9. *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 6.3.10 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 6.3.10. *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 $C \cap D \subseteq A = A \smallsetminus B$. By assumption, $C \cap D$ is finite, so we can apply Lemma 6.3.3 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 6.3.11 then asserts that (C, D) is smaller than all but finitely many (A_i, B_i).

Lemma 6.3.11. *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 6.2.1. 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 6.1), which partition $(C \cap B_i) \setminus S_i$, and finally give some $I \in \mathbb{N}$ satisfying all these conditions.

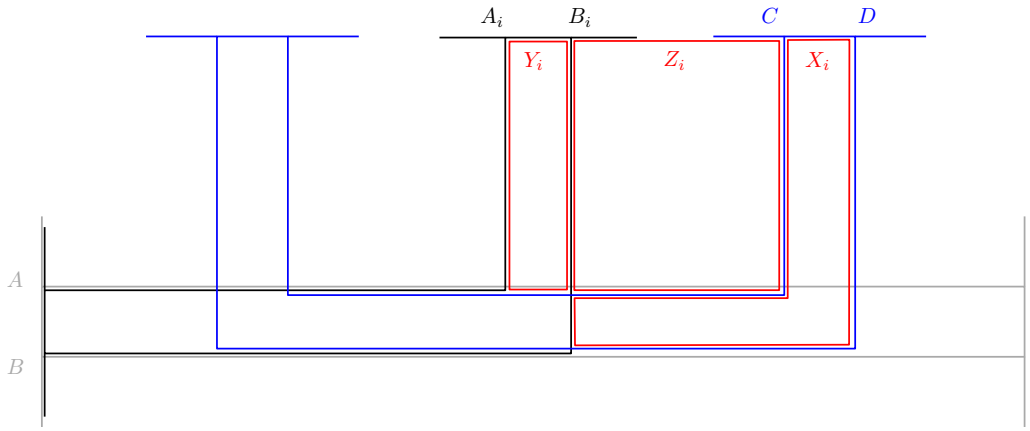

Figure 6.1: A diagram depicting the three sets X_i, Y_i and Z_i from the proof of Lemma 6.3.11.

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 6.3.3 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 6.3.7. So Lemma 6.3.8 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 6.3.3 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. □

6.4 The interlacing theorem

Interlaced strictly increasing sequences of separations have the same limit, as Lemma 6.3.2 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 6.3.2 and Lemma 6.3.5 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 [DJKK26, 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. [D17b, 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 $\overrightarrow{N}$ admits an interlaced sequence in $\overrightarrow{N}$ 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 6.4.1 (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 $\overrightarrow{N}$. 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 $\overrightarrow{N}$ 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 6.4.1.

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 6.4.1 is divided into the four Lemmas 6.4.3 to 6.4.6; 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 6.4.2. *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 6.4.3 shows that for the proof of Theorem 6.4.1, we can replace (IM2) with a weaker property (IM2') in terms of consecutive pre-tangles.

Lemma 6.4.3. *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 6.4.4 asserts, in the setting of Theorem 6.4.1, 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 6.4.3) – 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 6.4.4. *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 6.4.2 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 6.4.5 asserts, in the setting of Theorem 6.4.1, 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 6.4.5. *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 $(\vec{s}_i{}')_{i\in\mathbb{N}}$ 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 $(\vec{s}_i{}')_{i\in\mathbb{N}}$ in $\overrightarrow{N}$ which is strongly $\mathcal{P}$-relevant such that for every $k \in \mathbb{N}$, we have

(i) $\vec{s}_{f(k)}{}' = \vec{s}_k$, and

(ii) the pair $\vec{s}_i{}' < \vec{s}_{i+1}{}'$ is strongly $\mathcal{P}$-relevant for all $i < f(k)$.

Clearly, such a sequence $(\vec{s}_i{}')_{i\in\mathbb{N}}$ is as desired.

To start the recursion, we set $f(0) := 0$ and $\vec{s}_0{}' := \vec{s}_0$. Let $\ell \in \mathbb{N}$ such that we have already defined $f(0), \dots, f(\ell)$ and $(\vec{s}_i{}')_{i\leqslant f(\ell)}$ which satisfy properties (i) and (ii) for every $k \leqslant \ell$. We now proceed by defining $f(\ell+1) > f(\ell)$ and $\vec{s}_i{}'$ for the remaining $i = f(\ell)+1, \dots, f(\ell+1)$ such that $(\vec{s}_i{}')_{i\leqslant f(\ell+1)}$ 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 $(\vec{s}_i{}')_{i\leqslant f(\ell)}$ 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 $P_\ell, P_{\ell+1}, Q_{\ell+1} \in \mathcal{P}$ 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 $\vec{s}_{f(\ell+1)}{}' := \vec{s}_{\ell+1}$. Then $(\vec{s}_i{}')_{i \leqslant f(\ell+1)}$ 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 $f(\ell+1) := f(\ell)+2$, $\vec{s}_{f(\ell)+1}{}' := \vec{r}$ and $\vec{s}_{f(\ell)+2}{}' := \vec{s}_{\ell+1}$ yields the desired $(\vec{s}_i{}')_{i \leqslant f(\ell+1)}$ 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 6.4.2, 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 6.4.6 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 6.4.6. *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 6.4.3. 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 6.4.3 to 6.4.6 at hand, we are now ready to prove Theorem 6.4.1.

Proof of Theorem 6.4.1. Since $((A_i, B_i))_{i\in\mathbb{N}}$ is non-exhaustive, Lemmas 6.3.2 and 6.3.5 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 6.4.5 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 6.4.4 then yields a corresponding sequence $(P_i)_{i\in\mathbb{N}}$ in $\mathcal{P}$ such that (IM1) and (IM2') hold. Lemma 6.4.3 finally implies that (IM2) holds as well.

The 'moreover'-part can be obtained by applying Lemma 6.4.6 to $((A'_i, B'_i))_{i\in\mathbb{N}}$ and $(P_i)_{i\in\mathbb{N}}$. In order to apply Lemma 6.4.6, it remains to check that $(|A'_i, B'_i|)_{i\in\mathbb{N}}$ cofinitely exceeds every integer. By Lemma 6.3.2, 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 6.3.2 and Lemma 6.3.5 imply that $(|A'_i, B'_i|)_{i\in\mathbb{N}}$ cofinitely exceed every integer, as desired. □

We remark that Theorem 6.4.1 and its proof are purely about sequences of separations and do not refer to the underlying graph in any way. Thus, Theorem 6.4.1 can also be transferred to the more general framework of abstract separation systems which are extensively described in [D18a]. 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.

6.5 Proofs of Theorems 16 and 17

In this section, we prove our main results, Theorem 16 and its structural version Theorem 17, from which we will deduce Theorem 16. To prove Theorem 17, we will apply the interlacing theorem, Theorem 6.4.1, 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 [D17a].

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 6.5.1. *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 $K \in \mathbb{N}$ and 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 $((A_i, B_i))_{i\in\mathbb{N}}$ is exhaustive and all the $G[B_i]$ are connected, Lemma 6.3.10 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 [D17b, Proposition 8.2.2] implies

the following simple observation about combs and infinite sets of vertices in locally finite graphs.

Lemma 6.5.2. *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 6.3.6. Then Lemma 6.5.2 implies that some end of the graph is in the closure of the separator of such a 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 6.5.3. *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 6.3.6. So at least one end of G is in the closure of $A \cap B$ by Lemma 6.5.2. 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 6.3.9 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 $(C \smallsetminus D) \cap (A \cap B) \neq \emptyset$ as $C \cap D$ is finite. Arguing analogously for $(C, D) \in P_{\omega'}$, we obtain $(D \smallsetminus C) \cap (A \cap B) \neq \emptyset$. But then $\{A, B\}$ and $\{C, D\}$ cross by Lemma 6.2.1. □

We now use the interlacing theorem, Theorem 6.4.1, to show Theorem 6.5.4. 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 6.5.3 unique) end which is in the closure of the separator of its limit separation cannot be thin.

Theorem 6.5.4. *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.*

[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. [D17b, Section 8.6]).

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 6.4.1 to obtain a strictly increasing sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ in $\vec{N}$ 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 6.3.2 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 6.4.1, 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 6.5.3, 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. Every $\{A'_i, B'_i\}$ efficiently distinguishes P_i and P_ω.

Proof of the claim. 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 6.4.2 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 6.3.11, 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 6.3.9 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 6.2.1 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 6.5.1 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 6.5.4 by a contradiction in terms of cardinality via Lemma 6.5.1, 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 $G[B \setminus A]$ (see Section 6.6).

It remains to deduce our main results, Theorems 16 and 17, by applying Theorem 6.5.4 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 6.5.5 ([EKT22b, 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 17. Each separation in the tree of tangles N efficiently distinguishes some two tangles in G and is hence tight by Lemma 6.5.5. Since every tangle in G is also a pre-tangle in G, Theorem 6.5.4 implies that the unique end in the closure of $A \cap B$ is thick, as desired. □

Proof of Theorem 16. By Theorem 6.1.2, there exists a tree of tangles N of G. Since G does not have any thick end, Theorem 17 implies that every strictly increasing sequence in $\overrightarrow{N}$ is exhaustive, so N is itself exhaustive. Therefore, N induces a canonical tree-decomposition of G by Lemma 6.3.4, 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 6.1.2: Elbracht, Kneip and Teegen showed in [EKT22b, 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 16 and Theorem 17 for sets $\mathcal{P}$ of profiles with the same proofs.

6.6 Infinitely many rays beyond the limit

In this section we prove a structurally refined version of Theorem 6.5.4: while Theorem 6.5.4 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 6.6.1 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 6.6.1. *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 $\overrightarrow{N}$, 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 [H65, Satz 1] by Halin one easily obtains the following (see also [D17b, Theorem 8.2.5 & Exercise 43]).

Theorem 6.6.2. *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 6.6.1. We first proceed as in the proof of Theorem 6.5.4 to obtain a strictly increasing sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ in $\overrightarrow{N}$ 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 6.6.2, 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 6.3.3, 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 $\mathcal{R}_i$ 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 $\mathcal{R}_i$. For every $i \geqslant I$, we fix the subfamily $\mathcal{R}'_i$ of $\mathcal{R}_i$ 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 $\mathcal{R}'_i$ 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 $\mathcal{R}'_i$.

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 [K36, VI. Kapitel § 2, Satz 6] (see also [D17b, 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 $R_i^{\ell-1}$ is the subpath of R_i^ℓ of length $\ell-1$ starting in Z for $i = 1, \dots, k$. Let R_i be the ray $\bigcup_{\ell\in\mathbb{N}} R_i^\ell$ for $i = 1, \dots, k$. These rays R_i are disjoint as the R_i^ℓ 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 $\mathcal{R}_i$ for infinitely many $i \geqslant I$. All the rays in $\mathcal{R}_i$ avoid X_i except in their startvertices. Hence, every ray in $\mathcal{R}_i$ 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 6.3.3.

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 6.6.1 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 [EKT22b, Example 4.9].

Example 6.6.3. We construct a graph G (see Figure 6.2) 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_i^{n+1} with v_i^n 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 $w^m := u_1^{m+1} = v_1^m$. 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

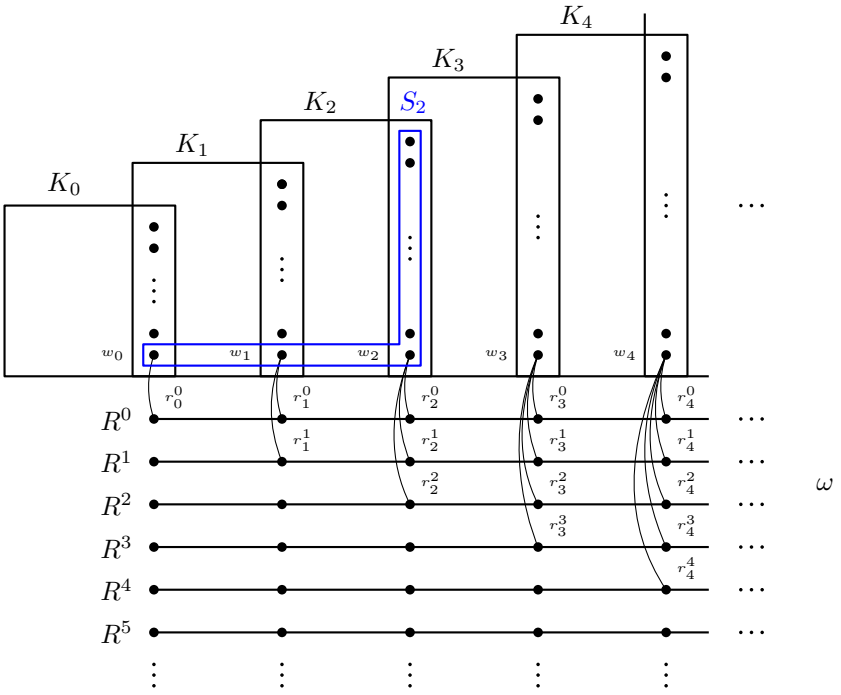

Figure 6.2: The graph G from Example 6.6.3

separation s_n with separator $S_n := \{v_1^i \mid i < n\} \cup \{v_j^n \mid j \leqslant 2^{n+1}\}$ 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.

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 $\overrightarrow{N}$, 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 6.6.3 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.

Acknowledgements

We thank Jan Kurkofka for an insightful discussion concerning the proof of Theorem 6.3.6 and Nathan Bowler for his idea how to prove Proposition 6.6.1. We also thank the reviewers for their valuable comments that greatly improved the presentation of some arguments and fixed a few mistakes. Especially, we are grateful for Reviewer 1's shorter and more intuitive proof of Lemma 6.3.9 and Reviewer 2's input regarding the presentation of Section 6.6.

PART III

Connectivity and tree-structures

CHAPTER 7

Linked tree-decompositions into finite parts

This chapter is based on [AJKP24] and joint work with Sandra Albrechtsen, Paul Knappe and Max Pitz.

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 G which admits a tree-decomposition into finite parts has a tree-decomposition into finite parts that displays all the ends of G and their combined degrees, resolving a question of Halin from 1977. This latter tree-decomposition yields immediate, unified proofs of the characterisations due to Robertson, Seymour and Thomas of graphs without half-grid minor, and of graphs without binary tree subdivision.

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 [T89] 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 [T90], 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

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

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 [T89],
- ◇ *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 18. *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 18 is additionally tight and componental.

Note that achieving just a subset of the properties in Theorem 18 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 [D94], 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.

[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 [D17b; RST95] done.

Having all three properties simultaneously is more challenging to achieve. In what follows, we hope to convince the reader of the usefulness of Theorem 18 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 [AJKP25] we provide examples showing that Theorem 18 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 [AJKP25, 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 18 displays the combinatorial and topological end structure of the underlying graph as follows, resolving a question by Halin from 1977 [H77, §6]. (Halin's original work yields proofs of Theorems 18 and 19 for locally finite connected graphs with at most two ends.)

Theorem 19. *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 [C19],
- ◇ *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).

[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 $\liminf_{e \in E(R_\varepsilon)} V_e = \bigcup_{n \in \mathbb{N}} \bigcap_{i \geqslant n} V_{e_i}$.

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 [D94] 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 19 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 [H78; P20], 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 19.

By Halin's grid theorem [H65; 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 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 19. 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 19. 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] [J69, §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 18 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.

[4] See [RST95; ST93] for details about the 'easy' implications.

[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.

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 [E18] for a modern treatment of the finite case.

Theorem 7.1.4 (Thomas 1990 [T90], 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 $e \in t_1Tt_2$ with $|V_e| < \ell$. There are two ways in which 'lean' is stronger than 'linked': First, every lean tree-decomposition $(T, \mathcal{V})$ satisfies that for *every* two nodes $s \neq t \in T$ there are $\min\{|V_e| \colon e \in E(sTt)\}$ pairwise disjoint V_s–V_t paths in G. In a way, this is the unrooted version of the property 'linked' as introduced above which only required this property when s and t are comparable. The other difference between 'linked' and 'lean' is that a bag of a lean tree-decomposition is only large if it is highly connected, which follows by considering the case $t_1 = t_2$ in the definition of 'lean'.

Now it is natural to ask whether Theorem 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 [AJKP25, 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 18 using the finite case of Theorem 7.1.4.

Theorem 20. *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 20 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 20 (see Theorem 20' 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 18 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 [P96] 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 18 we obtain a similar tree-decomposition which 'displays all the infinities' of the underlying graph: A rooted tree-decomposition $(T, \mathcal{V})$

- ◇ *displays the critical vertex sets* if the map $t \mapsto V_t$ restricted to the infinite-degree nodes of T whose V_t is finite is a bijection to the critical vertex sets. In this context, we denote for every critical vertex set X the unique infinite-degree node of T whose bag is X by t_X.

- ◇ *displays the tight components of every critical vertex set cofinitely* if it displays the critical vertex sets such that for every critical vertex set X every $G\mathring{\uparrow}e$ with $e = t_X t \in T$ with $t_X <_T t$ is a tight component of $G - X$ and cofinitely many tight components of $G - X$ are some such $G\mathring{\uparrow}e$, and

- ◇ *displays the infinities* if it displays the ends homeomorphically, their combined degree, their dominating vertices, the critical vertex sets and their tight components cofinitely.

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 21. *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 [C19, 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 [AJKP25, Construction 3.1] we present a planar graph witnessing that this result cannot be extended to graphs without $K^{\aleph_0}$ minor [AJKP25, 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 22. *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 22 as an application of Theorem 20: Our proof of Theorem 20 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 22.

An equivalent way of stating Corollary 22 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; EKT22b; EK22; JK24] extending the 'tree-of-tangles' theorem from Robertson and Seymour [RS91, (10.3)] to infinite graphs.

Additionally, the first author [A24b, Theorem 3] applies Theorem 21 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 [T89, (10.3)]. In light of the observation that all known counterexamples to well-quasi-orderings of infinite graphs [DL01; K95; P23; T88] do not have finite tree-width by the characterisation of finite tree-width graphs (i.e. graphs with normal spanning trees, Theorem 7.2.2 below) in [P21], 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 18 already satisfies Theorem 19. Section 7.4 consists of three parts. First, we give in Section 7.4.1 a brief overview over the proof of Theorem 18, which includes the statement of Theorems 23 and 24, our two main ingredients to the proof of Theorem 18. In Section 7.4.2, we state detailed versions of Theorems 18 and 21 which assert further properties of the respective tree-decompositions. In Section 7.4.3, we then reduce Theorems 18 and 21 to Theorems 23 and 24. We prove Theorem 23 in Section 7.5. In Section 7.6, we show some basic behaviour of paths and rays in torsos to prove Theorem 24 in Section 7.7. Finally, in Section 7.8, we prove Theorem 20 and Corollary 22. Furthermore, we compare in Section 7.A our Theorem 23 to Elm and Kurkofka's [EK22, Theorem 2].

[7] See Section 7.9 for definitions.

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

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 the chapter, graphs may be infinite. The notation and definitions follow [D17b] 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 $\mathrm{int}(\sigma) := \bigcap_{(A,B)\in\sigma} B$.

For a star σ of separations of G, we let $\mathrm{torso}(\sigma)$ be the graph arising from $G[\mathrm{int}(\sigma)]$ by making each separator of a separation in σ complete. We call these added edges that lie in $\mathrm{torso}(\sigma)$ but not in G *torso edges*. We refer to $\mathrm{torso}(\sigma)$ 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 $\mathrm{root}(T')$ is the node of T' containing $\mathrm{root}(T)$.

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 [D17b, 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 $\mathrm{torso}(\sigma_t)$ 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;

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

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* [J69, 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.*

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

$$\deg(\varepsilon) := \sup\{|\mathcal{R}| \mid \mathcal{R} \text{ is a family of disjoint rays in } \varepsilon\}.$$

Halin [H65, Satz 1] showed that this supremum is always attained. Together with the number $\operatorname{dom}(\varepsilon) := |\operatorname{Dom}(\varepsilon)|$ of vertices dominating ε, this sums up to the *combined degree* $\Delta_G(\varepsilon) := \deg(\varepsilon) + \operatorname{dom}(\varepsilon)$ of ε. Every end ε of a graph with normal spanning tree has countable combined degree: By [D17b, Lemma 8.2.3] $\deg(\varepsilon)$ is countable, and by [D17b, Lemma 1.5.5], every vertex in $\operatorname{Dom}(\varepsilon)$ lies on the unique normal ray in ε, so $\operatorname{dom}(\varepsilon)$ 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 $\operatorname{root}(T)$; we will frequently use this one-to-one correspondence between the ends of T and the rooted rays in T. Let $(T, \mathcal{V})$ be a rooted tree-decomposition $(T, \mathcal{V})$ of a graph G with finite adhesion, an end ε of G *lives in* an end η of T if some, or equivalently every, ray in ϵ has a tail in $G\mathring{\uparrow}e$ for every edge e of the unique rooted ray R in η; with the above relation between the rooted rays in T and the ends of T in mind, we also say that ε *gives rise to* R and R *arises from* ε. We remark that every end of G gives rise to at most one rooted ray in T, since $(T, \mathcal{V})$ has finite adhesion. If every end of G gives rise to some rooted ray of T, we encode this correspondence between the ends of G and T by a map $\varphi\colon \Omega(G) \to \Omega(T)$. In particular, this is the case if every torso of $(T, \mathcal{V})$ is rayless. If φ is a bijection between the ends of G and the ends of T, then we say that the rooted tree-decomposition $(T, \mathcal{V})$ *displays all ends* of G. If φ is a homeomorphism from $\Omega(G)$ to $\Omega(T)$, then $(T, \mathcal{V})$ *homeomorphically displays all ends* of G.

For a sequence $(V_i)_{i\in\mathbb{N}}$ of sets, we set $\liminf_{i\in\mathbb{N}} V_i := \bigcap_{i\in\mathbb{N}} \bigcup_{j\geqslant i} V_i$. If the sequence of sets V_{e_i} is indexed by the edges of a ray $R = v_0e_0v_1e_1\dots$ 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 $\liminf_{e\in R} V_e = \operatorname{Dom}(\varepsilon)$ for every rooted ray R of T and its arising end ε of G, then $(T, \mathcal{V})$ *displays the dominating vertices of every end*.

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

Lemma 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 $\mathcal{C}_X := \mathcal{C}(G - X)$ for the set of components of $G - X$ and $\breve{\mathcal{C}}_X := \breve{\mathcal{C}}(G - X) \subseteq \mathcal{C}_X$ 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 $\breve{\mathcal{C}}_X$ 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 [P96, 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 [D17b, 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 20 and 21, 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* $X \in \operatorname{crit}(\operatorname{torso}(\sigma))$*, then there are infinitely many tight components of* $G - X$ *which meet* $\operatorname{torso}(\sigma)$*.*

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

Proof. We remark that (i) immediately yields that $\operatorname{crit}(\operatorname{torso}(\sigma)) \subseteq \operatorname{crit}(G)$. 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 $\operatorname{torso}(\sigma) - X$ meet U. For every component C' of $\operatorname{torso}(\sigma) - X$ 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 $\operatorname{torso}(\sigma)$, 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 $\operatorname{torso}(\sigma) - X$ than C'. It thus suffices to show for all infinitely many components C' of $\operatorname{torso}(\sigma) - X$ 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 $X \subseteq \operatorname{int}(\sigma)$. It suffices to show that $C' := C \cap \operatorname{int}(\sigma)$ is a tight component of $\operatorname{torso}(\sigma) - X$ for every tight component C of $G - X$ which meets $\operatorname{int}(\sigma)$ but avoids the finite set U. The definition of torso immediately yields that C' induces a connected subgraph of $\operatorname{torso}(\sigma) - X$ with $N_{\operatorname{torso}(\sigma)}(C') \supseteq X$ because $X \subseteq \operatorname{int}(\sigma)$. It remains to show that $N_{\operatorname{torso}(\sigma)}(C') \subseteq X$. 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 19 follows from Theorem 18, 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*

[10] Our definition of *componental* agrees with the definition of *upwards connected* from [KKP23].

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 $\liminf_{e \in R} V_e = \mathrm{Dom}(\varepsilon)$ for every rooted R of T and the unique end ε of G which gives rise to R.*

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

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

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

Since $\liminf_{e \in R} S_n = \mathrm{Dom}(\varepsilon)$, we may assume that $S_n \cap S_m \subseteq \mathrm{Dom}(\varepsilon)$ for all $m \geqslant n$ by passing to a further subsequence. We also have $C_G(S_m, \varepsilon) \subseteq C_G(S_n, \varepsilon)$ since ε gives rise to R. Moreover, $\bigcap_{n \in \mathbb{N}} C_G(S_n, \varepsilon) = \emptyset$, since ε gives rise to R and $(T, \mathcal{V})$ is a tree-decomposition. This implies by [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 $\mathcal{R} = \bigcup_{n \in \mathbb{N}} \mathcal{P}_n$ is a family of pairwise disjoint paths and rays; the rays in $\mathcal{R}$ are necessarily ε-rays since ε is the unique end of G that gives rise to R, and the finite paths in $\mathcal{R}$ necessarily have a vertex of $\mathrm{Dom}(\varepsilon)$ as their last vertex. Hence, $|\mathcal{R}| \leqslant \Delta(\varepsilon)$. Moreover, we clearly have $|\mathcal{R}| = \liminf_{n \in \mathbb{N}} |S_n|$, which implies that $\liminf_{n \in \mathbb{N}} |S_n| = |\mathcal{R}| \leqslant \Delta_G(\varepsilon)$, as desired. □

Proof of Theorem 19 given Theorem 18. By Lemmas 7.3.1 to 7.3.3 the tree-decomposition from Theorem 18 is as desired. □

7.4 A high-level proof of the main result of this chapter

7.4.1 A two-step approach to Theorem 18

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 23. *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 24. *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 23 and Theorem 24 can be combined to yield a proof of Theorem 18. The subsequent Sections 7.5 and 7.7 are then concerned with proving Theorems 23 and 24 respectively.

7.4.2 Detailed versions of the main theorems

In fact, we will prove a version of Theorem 18 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 18' (Detailed version of Theorem 18)**.** *Every graph G of finite tree-width admits a rooted tree-decomposition $(T, \mathcal{V})$ into finite parts that is linked, tight, and componental. In particular, $(T, \mathcal{V})$ displays all the ends of G homeomorphically, their combined degrees and their dominating vertices. Moreover, we may assume that*

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

[11] This is not completely true: In order to ensure that the arising tree-decomposition is linked we need a 'refinement version' of Theorem 24 (see Theorem 24') 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 24.

(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 $G{\restriction}e'$, or is contained in a critical vertex set of G that is included in $G{\restriction}e'$, and*

(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 [H75, 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; [AJKP25, 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 21 and then derive Theorem 18 from it. In fact, we derive the detailed version Theorem 18' from the following detailed version of Theorem 21.

Theorem 21' (Detailed version of Theorem 21)**.** *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 18'. Moreover,*

(I1) *if $G\mathring{\restriction}e$ 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 21' 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 $V_t \in \operatorname{crit}(G)$ 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 18' given Theorem 21'. Let $(T', \mathcal{V}')$ be the tree-decomposition obtained from Theorem 21'. 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 18', 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 $X \in \operatorname{crit}(G)$ can appear as at most two bags of $(T', \mathcal{V}')$, which then need to be adjacent. So assume $V_s = V_t = X$ with s being the unique down-neighbour of t. Then by (I2), $\deg(t) = \infty$. Since $(T', \mathcal{V}')$ displays the critical vertex sets and their tight components cofinitely, it follows that $G\mathring{\uparrow}(st)$ is disconnected. Thus, we have contracted the edge st in the construction of $(T, \mathcal{V})$. Hence, also every critical vertex set of G appears as at most one bag of $(T, \mathcal{V})$, so its bags are paiwise distinct. □

In order to prove Theorem 21', we still follow the promised two-step approach, but need the following detailed versions of Theorems 23 and 24.

Theorem 23' (Detailed version of Theorem 23)**.** *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 21'.*

In order to state the detailed version of Theorem 24, 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 24' (Detailed version of Theorem 24)**.** *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 $s \neq r := \mathrm{root}(T)$. 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 $e = rt \in T$.*

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

7.4.3 Proof of the main result

According to our two-step approach, we prove Theorem 21' by applying Theorem 24' to the torsos of the tree-decomposition given by Theorem 23'. 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 $V'_t := V_t \cap \mathrm{int}(\sigma)$ is a tree-decomposition of the torso at σ in G, as desired. □

Proof of Theorem 21' given Theorem 23' and Theorem 24'. Let $(T^1, \mathcal{V}^1)$ be the rooted tree-decomposition from Theorem 23' 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 21'. Let t be a node of T^1. We describe how we refine the torso at t in T^1 using Theorem 24':

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 24'.

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 24'. 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 24' 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 21', since $(T^1, \mathcal{V}^1)$ and all the $(T^t, \mathcal{V}^t)$ satisfy (I1). It also satisfies (I2). Indeed, by Theorem 23', $(T^1, \mathcal{V}^1)$ satisfies (I2). Moreover, by (R4) from Theorem 24', 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 24' satisfies $X_t = V^1_{st} \subsetneq V^1_t = \mathrm{int}(\sigma) = \mathrm{int}(\sigma'_t)$ by (I2) from Theorem 23' 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 24' 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 23', which we restate here for convenience.

Theorem 23' (Detailed version of Theorem 23)**.** *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 21'.

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 23' 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 $X \in \operatorname{crit}(G)$ 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 $\operatorname{torso}(\sigma_t)$ is not tough, that is, there is some $X \in \operatorname{crit}(\operatorname{torso}(\sigma_t))$. 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 23' 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 $X \in \mathrm{crit}(G)$, let t'_X be the node of T' whose branch set in T contains t_X. Since we only contracted edges of T whose adhesion set is not a critical vertex set of G, we still have that cofinitely many tight components of $G - X$ are of the form $G\mathring{\uparrow}e$ for an edge $e = t'_X s \in T'$ with $t'_X <_{T'} s$. Hence, we may apply Lemma 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 23' 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 $X \in \mathrm{crit}(G)$ cofinitely many tight components of $G - X$ are of the form $G\mathring{\uparrow}e$ for an edge $e = t'_X t \in T'$ with $t'_X <_{T'} t$. 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 23'. By applying Construction 7.5.4 to the tree-decomposition $(T, \mathcal{V})$ from Lemma 7.5.3, we obtain a tree-decomposition$(T', \mathcal{V}')$. By Lemma 7.5.6 $(T', \mathcal{V}')$ 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 23' was proven by Elm and Kurkofka [EK22, Theorem 2]. However, their result cannot be used in place our Theorem 23', 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.A for a detailed discussion of the differences between the two results.

7.6 Lifting paths and rays from torso

From Theorem 23' we obtain a tree-decomposition $(T, \mathcal{V})$ whose torsos are tough. Following our overall proof strategy for Theorem 18 we want to further decompose each of its torsos $\mathrm{torso}(\sigma_t)$ 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 $\mathrm{torso}(\sigma_t)$ 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 $A_e \smallsetminus B_e$. 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 $A \smallsetminus B$ 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. [D17b, 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 24', which we restate here for convenience.

Theorem 24' (Detailed version of Theorem 24)**.** *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 $G{\uparrow}e$ or linked to a set $A \cap B \subseteq G{\uparrow}e$ with $(A, B) \in \sigma$,*

(R3) *for every $e <_T e' \in E(T)$ with $|V_e| \leqslant |V_{e'}|$, each vertex of $V_e \cap V_{e'}$ either dominates some end of G that lives in $G{\uparrow}e'$ or is contained in $A \cap B \subseteq V(G{\uparrow}e')$ for some $(A, B) \in \sigma$,*

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

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

We remark that in Theorem 24' 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 24' to the following statement which differs slightly from Theorem 24'. 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 $X \subseteq \mathrm{int}(\sigma)$ be a finite set of vertices of G. Assume further that for every $(A, B) \in \sigma$, the graph $G[A \cap B]$ is complete and $A \cap B \subseteq \mathrm{Dom}(G[A \setminus B])$.*

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

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

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

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

(R4') $(T, \mathcal{V})$ *satisfies* (R4) *from Theorem 24'.*

Proof of Theorem 24' given Theorem 7.7.1. Let H be the graph obtained from torso(σ) 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 24', torso(σ) 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 $V_e = V'_e \subseteq \text{int}(\sigma)$ is linked to $A \cap B \subseteq V(G{\restriction}e)$ since $A \cap B$ separates int(σ) 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 $G{\restriction}e$. Indeed, let $\mathcal{R}$ be a family of $|V_e|$ equivalent, disjoint rays in torso(σ) that start in V_e and that witness that V_e is end-linked. Then Lemma 7.6.3 (i) yields a family $\mathcal{R}'$ of $|V_e|$ disjoint rays in $G{\restriction}e$ that start in V_e and that each meet $V(\mathcal{R})$ infinitely often. In particular, since $\mathcal{R}$ is finite, for every $R' \in \mathcal{R}'$ there is $R \in \mathcal{R}$ such that R' meets $V(R)$ infinitely often. To see that the rays in $\mathcal{R}'$ are equivalent, let $R'_0, R'_1 \in \mathcal{R}'$ be given. Since R_0 and R_1 are equivalent in torso(σ), 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 torso(σ). Since these paths are in fact R_0'–R_1' paths in torso(σ), 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 $Y \subseteq V(G{\restriction}e)$ such that the adhesion set V_e^n is contained in Y. We then replace the bag V_ℓ^n with Y and add for each component C of $(G{\restriction}e) - Y$ 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 $G\mathring{\restriction}e$ 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 $\mathcal{D}$ of pairwise non-touching end-linked $(<\aleph_0)$-regions D that are disjoint from X and satisfy $X \cap N_H(D) \subseteq \mathrm{Dom}(D)$; a vertex $x \in V(G) \setminus X$ that lies in no $D \in \mathcal{D}$.

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

Recursion: Iterate the following step:

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

(N1) C_i is such an ℓ_i-region where $\ell_i \in \mathbb{N}$ is minimum among such regions, and

(N2) C_i is an inclusion-wise maximal such region subject to (N1).

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

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

Recall that we will construct the tree-decomposition for Theorem 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

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

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

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

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

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

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

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

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

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

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

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

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 $\mathcal{D}$ 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 [D17b, 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 $\mathcal{D}$ and with the C_j for $j < i$. For this, we first need one auxiliary lemma.

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

Lemma 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 [D17b, Proposition 8.4.1]), there is an X–Y separator S of size less than $\min\{|X|, |Y|\}$ in $H[V(C) \cup X \cup Y]$. Since C is end-linked, there is some end ε of H which lives in C and a family $\{R_v \mid v \in N(C)\}$ of pairwise disjoint $N(C)$–ε paths and rays with $v \in R_v$. Since $|S| < \min\{|X|, |Y|\}$ and the R_v are pairwise disjoint, there is $x \in X$ and $y \in Y$ such that R_x and R_y both avoid S. Since R_x and R_y are $N(C)$–ε paths or rays, there is some R_x–R_y path P in C which avoids the finite set S. Hence, $R_x + P + R_y$ is a connected subgraph of $H[V(C) \cup X \cup Y]$ which meets X and Y but avoids S. This contradicts the fact that S is an X–Y separator in $H[V(C) \cup X \cup Y]$. □

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

Lemma 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, $D \in \mathcal{D}$ meets $N(C)$.

Let $\mathcal{E}'_<$ consist of all $D \in \mathcal{D}'$ with $|V(D) \cap N(C)| < |V(C) \cap N(D)|$. Now $C^* := G[V(C) \cup \bigcup_{D \in \mathcal{E}'_<} V(D)]$ has neighbourhood $(N(C) \setminus \bigcup_{D \in \mathcal{E}'_<} V(D)) \cup \bigcup_{D \in \mathcal{E}'_<} (N(D) \setminus V(\bar{C}))$, 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

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

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

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

Lemma 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 $D \in \mathcal{D}$ 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 $D \in \mathcal{D}$ is contained in a component of $H - Y$.*

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

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

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

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

(A6) *Set $H^0 := H$ and $Y^0 := Y$. Let $C'_1 \supseteq C'_2 \supseteq \cdots \supseteq C'_n$ 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 $\mathcal{D}^m := \{D \in \mathcal{D} \mid D \subseteq C'_m\}$ and an arbitrary vertex $x_m \in C'_m - (\bigcup \mathcal{D})$. Then there are $\min\{|N_H(C'_m)| \mid m \in \{1, \dots, n\}\}$ pairwise disjoint X–$N_H(C'_n)$ paths in H.*

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

Proof. (A1): Assume that $x \notin Y$. Let $\mathcal{C}_x$ be the collection of all those C_i that contain x. Then $\mathcal{C}_x$ is non-empty since by the definition of Y at least one C_i contains x. We claim that $C := \bigcup \mathcal{C}_x$ is a component of $H - Y$ with $|N(C)| < k = |X|$. Indeed, it is immediate from Algorithm 7.7.2 that every $C_i \in \mathcal{C}_x$ is a region as in Case A, since the regions that where chosen in Case B avoid x; in particular, every $C_i \in \mathcal{C}_x$ is a $(< k)$-region. Since all the $C_i \in \mathcal{C}_x$ meet in x and are nested, $\mathcal{C}_x$ is a chain (with respect to inclusion). Thus, Lemma 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 $C = \bigcup \mathcal{C}_x$, we have $N(C) \subseteq \bigcup\{N(C_i) \mid C_i \in \mathcal{C}_x\}$. Since all the C_i are nested, every C_j which meets the neighbourhood of some $C_i \in \mathcal{C}_x$ already contains C_i; in particular, such C_j are also in $\mathcal{C}_x$ as $x \in C_i \subseteq C_j$. Thus $N(C)$ meets no C_i, and so the definition of Y yields that $N(C) \subseteq Y$.

(A2): By the assumptions on $\mathcal{D}$ and the choice of the C_i, in every step i of Algorithm 7.7.2 every $D \in \mathcal{D}$ is a candidate for the region C_i in Case B as long as D is ε-linked for some end ε of H which lives in no C_j for $j < i$. Again by the assumptions on $\mathcal{D}$, the region $D \in \mathcal{D}$ is ε-linked for some end ε of H. So, since Algorithm 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 $\mathcal{D}$ that in step i, the region $D \in \mathcal{D}$ is nested with all C_j for $j < i$, disjoint from X and that $X \cap N(D) \subseteq \mathrm{Dom}(D)$. Since D is ε-linked and ε lives in $C_i \subseteq D$, there are $|N(D)|$ disjoint $N(D)$–$N(C_i)$ paths in $\bar{D} - C_i$. These paths can be extended to disjoint $N(D)$–ε' paths and rays where ε' is the end so that C_i is ε'-linked. In particular, D is end-linked to the end ε' which lives in no C_j for $j < i$, and $\ell_i = |N(C_i)| \geqslant |N(D)|$. Thus, D was a candidate for C_i and satisfies (N1), as desired.

(A3): Suppose towards a contradiction that there is an end ε of H that lives in no C_i. For every $\ell \in \mathbb{N}$, let $\mathcal{C}^{<\ell}_{\max}$ 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 $C_j \in \mathcal{C}^{<\ell}_{\max}$. Moreover, since all C_i are nested, the regions in $\mathcal{C}^{<\ell}_{\max}$ 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 $\mathcal{D} \cup \mathcal{C}^{<k}_{\max}$, 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 $\mathcal{D} \cup \mathcal{C}^{<k}_{\max}$. Set $\ell := |N(C)|$. Then Lemma 7.7.6 applied to C and the set $\mathcal{E}$ of all $\subseteq$-maximal regions in $\mathcal{D} \cup \mathcal{C}^{<\ell+1}_{\max}$ yields a ε-linked $(\leqslant \ell)$-region C' which is nested with all regions in $\mathcal{D} \cup \mathcal{C}^{<\ell+1}_{\max}$ 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 $D \in \mathcal{C}^{<\ell+1}_{\max} \setminus \mathcal{C}^{<k}_{\max}$ 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 $\mathcal{D} \cup \mathcal{C}^{<\ell+1}_{\max}$. In particular, C' is nested with all $D \in \mathcal{D}$ 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 [D17b, 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 [D17b, 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^N_i 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 $\mathcal{C}^{\leqslant|S|}_{\max}$ be the $\subseteq$-maximal elements of the C^N_i with $\ell^N_i \leqslant |S|$. Note that by Observation 7.7.8 (O5) every C^N_i with $\ell^N_i < |S|$ is contained in some $C^N_j \in \mathcal{C}^{\leqslant|S|}_{\max}$. Since all regions in $\mathcal{C}^{\leqslant|S|}_{\max} \cup \mathcal{D}$ are end-linked, they are well linked by Lemma 7.7.5. Moreover, if ε lives in $D \in \mathcal{D}$, then $D \subseteq C'_n \subseteq C$ by (A2). To be able to apply Lemma 7.7.6 to C and the set $\mathcal{E}$ of $\subseteq$-maximal regions in $\mathcal{C}^{\leqslant|S|}_{\max} \cup \mathcal{D}$ it suffices that ε lives in no C^N_j with $\ell^N_j \leqslant |S|$.

So suppose for a contradiction that ε lives in some C^N_j with $\ell^N_j \leqslant |S|$. We claim that then $C'' := C^N_j \cup C'_n$ is an $(\leqslant |S|)$-region. Indeed, C'' is a region because C^N_j and C'_n are both connected and intersect as ε lives in both of them. Moreover, $N(C'') = (N(C^N_j) \cup N(C'_n)) \setminus V(C'') = (N(C^N_j) \setminus V(C'_n)) \dot\cup (N(C'_n) \setminus V(\bar{C}^N_j))$ has size at most $|N(C^N_j)| = \ell^N_j \leqslant |S|$: if not, one easily checks by doubling counting that $Z := ((N(C^N_j) \cup N(C'_n)) \cap V(C'')) \cup (N(C^N_j) \cap N(C'_n)) = (N(C'_n) \cap V(C^N_j)) \dot\cup (N(C^N_j) \cap V(C'_n)) \dot\cup (N(C^N_j) \cap N(C'_n))$ has size less than $|N(C'_n)|$. But this contradicts that C'_n is ε-linked as ε lives in a component of $C^N_j \cap C'_n$, whose neighbourhood then is contained in Z, and so all $N(C'_n)$–ε rays and paths have to meet Z. Hence, C'' is a $(\leqslant |S|)$-region with $C'_n \subseteq C''$. Moreover, C^N_j and C'_{N+1} both contain C'_n by assumption and thus touch. So since C^N_j and C'_{N+1} are nested by Algorithm 7.7.2 (because C'_{N+1} is some C^N_i by (A4)), Observation 7.7.8 (O4) implies that $C^N_j \subseteq C'_{N+1}$, and therefore $C'' \subseteq C'_{N+1}$. But this contradicts our choice of C. Thus, ε lives in no region in $\mathcal{C}^{\leqslant|S|}_{\max} \cup \mathcal{D}$.

So we may apply Lemma 7.7.6 to C and the set $\mathcal{E}$ of $\subseteq$-maximal separations in $\mathcal{C}^{\leqslant|S|}_{\max} \cup \mathcal{D}$ to obtain an ε-linked $(\leqslant|S|)$-region C' which is disjoint from X and which, for every region in $\mathcal{C}^{\leqslant|S|}_{\max} \cup \mathcal{D}$, either contains that region or does not touch it. Thus, C' is nested with all C^N_i with $\ell^N_i \leqslant |S|$ and with all $D \in \mathcal{D}$. This concludes the proof since the existence of C' then contradicts that C^N_I satisfies (N1) where $I = \min\{i \mid \ell^n_i > |S|\}$. Here we used that C'_{N+1} is some C^N_i 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 $\mathcal{D}$ the set of subgraphs $G[A \setminus B]$ of G with $(A, B) \in \sigma$. The assumptions on σ ensure that $\mathcal{D}$ satisfies the assumptions in Algorithm 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 $V^0_r := V(G)$. 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 $V^m_t = V^n_t$ for all $t \in T^m$ that are not at height m, the decomposition tree T^m has height m, all their leaves are on height m except possibly those whose corresponding leaf separation is a (B, A) with $(A, B) \in \sigma$,

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

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

Now apply Algorithm 7.7.2 to $H := G[V^n_\ell]$ with the finite set X', the vertex x and $\{D \in \mathcal{D} \mid D \subseteq H\}$ to obtain $Y \subseteq V(H)$. Then let T^ℓ be the star with centre ℓ and whose set of leaves is the set $\mathcal{C}$ of all components of $H - Y$. Further, set $V^\ell_\ell := Y$ and $V^\ell_C := \bar{C}$ for every $C \in \mathcal{C}$. Then Theorem 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 $|V^\ell_e| \geqslant |V^n_f|$ for an edge $e \in T^\ell$, then $V^\ell_e \cap V^n_f \subseteq \operatorname{Dom}(H\uparrow e)$.

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

follows immediately from the construction and the discussion of the properties of the $(T^\ell, \mathcal{V}^\ell)$. Its ($X$-)linkedness is ensured by Theorem 7.7.9 (A6).

Let $(T, \mathcal{V})$ be the limit of $(T^n, \mathcal{V}^n)$ for $n \to \infty$, that is $T = \bigcup_{n\in\mathbb{N}} T^n$, $\mathrm{root}(T) := r = \mathrm{root}(T^n)$ 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 $D \in \mathcal{D}$. 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 $v_i = \min_{\leqslant_S}(V(G\mathring{\uparrow}e^n) \setminus (\bigcup V(\mathcal{D})))$ 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 $D \in \mathcal{D}$. 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 $\mathcal{D}$ 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 20, the strengthening of Theorem 18 for graphs without half-grid minor in which we replace the linkedness of the tree-decomposition with leanness. We restate it here in its more detailed version:

Theorem 20' (Detailed version of Theorem 20)**.** *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 20' is structured as follows. We start our construction with the tree-decomposition $(T, \mathcal{V})$ of the graph G from Theorem 21', 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 20', 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 20' 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 [A24b, 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 20' 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 20', 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 [A24b], 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

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

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 ([A24b, 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* $\mathrm{torso}(\sigma)$ *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 [D17b, 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 20': the pre-processing step in which we contract certain edges of the tree-decomposition from Theorem 21. 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 $w = \mathrm{tw}(G) \in \mathbb{N}$ or $m = 0$ for graphs G whose tree-width is not finitely bounded. Analogously, we define *right-good*. We call (A, B) *good* if it is both left- and right-good. The bounds in the definition of (left-)good are exactly the ones sufficient to apply Lemmas 7.8.1 and 7.8.2 in the respective contexts.

Setting out from Theorem 21', 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 21'. 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 21' 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 20'. 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 [H78]. Let $(T, \mathcal{V})$ be the (rooted) tree-decomposition of G from Theorem 21', 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 [T90, 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 $(T^t, \mathcal{V}^t)$ is the trivial tree-decomposition consisting of a single node-tree, since the torso of $(T', \mathcal{V}')$ at t is complete.

[14] We remark that this is not explicitly stated in [T90, Theorem 5] but follows directly from its proof.

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 $V'_t \in \mathrm{crit}(G)$ 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 $\tilde{V}_{s_1} \in \mathrm{crit}(G)$, 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 $\tilde{V}_{s_2} \in \mathrm{crit}(G)$.

Suppose now that $\tilde{V}_{s_1}, \tilde{V}_{s_2} \notin \mathrm{crit}(G)$. Since there are no k disjoint Z_1–Z_2 paths in G, Menger's theorem (see for example [D17b, Proposition 8.4.1]) yields a separation (C, D) of G of order less than k with $Z_1 \subseteq C$ and $Z_2 \subseteq D$. We now use (C, D) to find a Z_1–$\tilde{V}_f$ or a Z_2–$\tilde{V}_f$-separator of size less than k, which gives the desired contradiction.

Let (A, B) be the separation of G induced by $\vec{f} = (s_1, s_2)$, and let $X := (A \cap B) \cap (C \setminus D)$ and $Y := (A \cap B) \cap (D \setminus C)$. By symmetry on the assumptions up to this point, we may assume $|Y| \leqslant |X|$. Since $\tilde{V}_{s_1} \notin \mathrm{crit}(G)$, (A, B) is left-good by the construction of $(\tilde{T}, \tilde{\mathcal{V}})$ and (i) from Lemma 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 $(C \cap D) \cap (A \smallsetminus B)$, so $|(C \cap D) \cap (A \smallsetminus B)| \geqslant |Y|$. But then

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

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 [D17a, Theorem 3], for every infinite tangle τ of a graph G there is either an end ε of G such that a finite-order separation (A, B) lies in τ if and only if B contains a tail of every, or equivalently some, ray in ε, or there is a critical vertex set X of G such that $(V(C) \cup X, V(G - C)) \in \tau$ for all $C \in \mathcal{C}(G - X)$. If the first case holds for τ, then we say that τ is *induced* by ε.

Following [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 $C \in \mathcal{C}_X$ and such that $(V(G - C), V(C) \cup X) \in \tau'$ for a component $C \in \mathcal{C}_X$.

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

Observation 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 [AJKP25, Construction 3.1] has finite tree-width and thus by Theorem 21 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 [AJKP25, 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 22' (Tangle version of Corollary 22)**.** *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 20' 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 $C \in \mathcal{C}(G - X_i)$, then there is unique infinite-degree node $t_i \in T$ with $V_{t_i} = X_i$, and we then let f_i be the unique edge of the $t_i T e$ path incident with t_i. Otherwise, τ_i is induced by an end ε_i of G, and we then let f_i be any edge on the unique ray R of T starting in e and arising from ε_i such that $|V_f| \geqslant |V_{f_i}|$ for all edges $f >_T f_i$ on R.

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

Proof of Corollary 22. Apply Corollary 22'. □

7.A A closer look at Theorem 23'

A theorem which is similar to Theorem 23' was proven by Elm and Kurkofka [EK22, Theorem 2]:

Theorem 7.A.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.*

In the remainder of this section we compare Theorem 23' and Theorem 7.A.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 23' immediately implies Theorem 7.A.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

[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 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.

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.A.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

$$\begin{aligned}&\{\{V(G) \setminus \bigcup \mathcal{K}(X), X \cup \bigcup \mathcal{K}(X)\} \mid X \in \operatorname{crit}(G)\}\\ \cup\ &\{\{V(G) \setminus C, X \cup C\} \mid X \in \operatorname{crit}(G), C \in \mathcal{K}(X)\}\end{aligned} \tag{7.1}$$

for some choice of $\mathcal{K}(X) \subseteq \breve{\mathcal{C}}(G - X)$ containing all but at most one element of $\breve{\mathcal{C}}(G - X)$.

The nested set induced by the tree-decomposition $(T', \mathcal{V}')$ from Theorem 23' is also of the form as in (7.1) but for some choice of $\mathcal{K}(X) \subseteq \mathcal{C}(G - X)$ consisting of cofinitely many (instead of one) elements of $\breve{\mathcal{C}}(G - X)$. 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 $\mathcal{K}(X)$ containing all but at most one element of $\mathcal{C}(G - X)$. In particular, we cannot strengthen the definition of 'displaying the tight components of all critical vertex sets cofinitely' by dropping the 'cofinitely'.

Example 7.A.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 $\mathcal{K}(X) \subseteq \breve{\mathcal{C}}(G - X)$ containing all but at most one element of $\breve{\mathcal{C}}(G - X)$.

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 $\mathcal{K}(X) \subseteq \breve{\mathcal{C}}(G - X)$ containing all but at most one element of $\breve{\mathcal{C}}(G - X)$. 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

$$\begin{aligned}C_i &:= G[\{r_j \mid j \in 2\mathbb{N} + 1, j \leqslant i - 2\} \cup \{u_{jk} \mid j < i, k \in \mathbb{N}\}] \text{ or}\\ C_i' &:= G[\{r_j \mid j \geqslant i\} \cup \{u_{jk} \mid j > i, k \in \mathbb{N}\} \cup \{d\}].\end{aligned}$$

In particular, for all X_i all components of $G - X_i$ are tight (cf. Figure 7.2). Hence, the

[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.

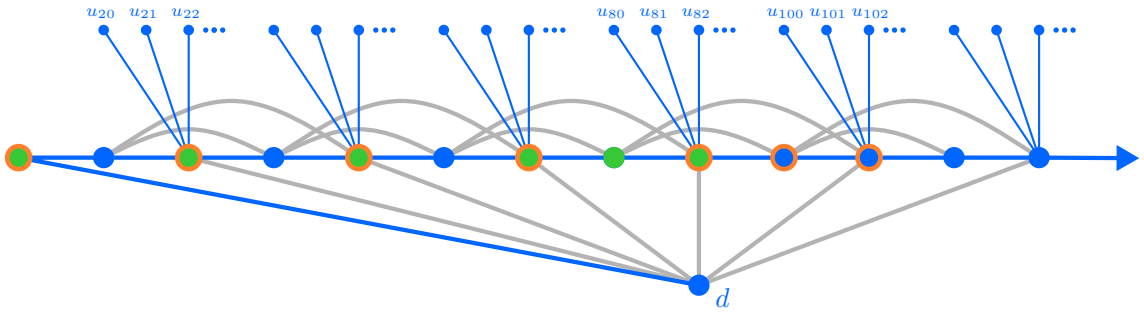

Figure 7.1: Depicted is the graph H. The green (orange) vertices are precisely the neighbours in G of u_{8j} (u_{10j}) for all $j \in \mathbb{N}$. The blue edges induce a normal spanning tree of G with root d.

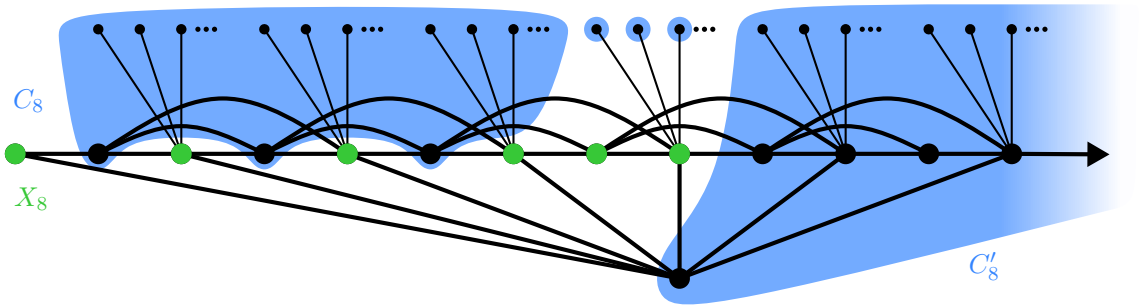

Figure 7.2: Indicated in green is the critical vertex set X_8 of G. The components of $G - X_8$ are indicated in blue.

unique nested set that is of the form as in (7.1) for a choice of $\mathcal{K}(X_i) \subseteq \breve{\mathcal{C}}(G - X_i)$ containing all but at most one element of $\breve{\mathcal{C}}(G - X_i)$ is the set

$$N := \{\{V(C) \cup X_i, V(G - C)\} \mid X_i \in \operatorname{crit}(G), C \in \mathcal{C}(G - X_i)\}.$$

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 $G\mathring{\uparrow}e_i = V(C'_i)$. 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 [AJKP25] and joint work with Sandra Albrechtsen, Paul Knappe and Max Pitz.

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.

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 [T89] 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 [T90], 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 25. *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 25 witnesses that the answers to both questions (i) and (ii) are in the negative.

On the positive side, we provide in [AJKP24, 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 25 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 25 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; E18]: 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 [T89] 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 [T89] 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. [AJKP24, §1.2-1.4] and [A24b]; 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 25:

Example 26. *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 $\kappa = \aleph_0$ [AJKP24, 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. [AJKP24, §1.2-1.4] and [A24b]): In [AJKP24, §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 [H77, §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 [AJKP24, §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 [AJKP24, §1.4] and [A24b].

In fact, we show in [AJKP24, 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 27. *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 25 and 26, and in Section 8.4 we prove Example 27. 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 [AJKP24, §2] for all the relevant definitions. Besides that, our notation and definitions follow [D17b]. 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 [C19],
- ◇ *displays all dominating vertices* if $\liminf_{e\in E(R_\varepsilon)} V_e = \mathrm{Dom}(\varepsilon)$[2] for all $\varepsilon \in \Omega(G)$, and

[1] We remark that this definition extends the one given in [AJKP24, §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 $\liminf_{e\in E(R_\varepsilon)} V_e = \bigcup_{n\in\mathbb{N}} \bigcap_{i\geqslant n} V_{e_i}$.

◇ *displays all combined degrees* if $\liminf_{e\in E(R_\varepsilon)} |V_e| = \Delta(\varepsilon)$ for all $\varepsilon \in \Omega(G)$.

Moreover, we recall the following lemma [AJKP24, 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 [AJKP24, 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 25 and 26 – negative answers to questions (i), (ii) and (iv)

In this section we explain Examples 25 and 26, which we restate here for convenience:

Example 25. *There is a planar, locally finite, connected graph that admits no lean tree-decomposition.*

Example 26. *There is a planar, locally finite, connected graph which admits no tight, componental rooted tree-decomposition into finite parts that is strongly linked.*

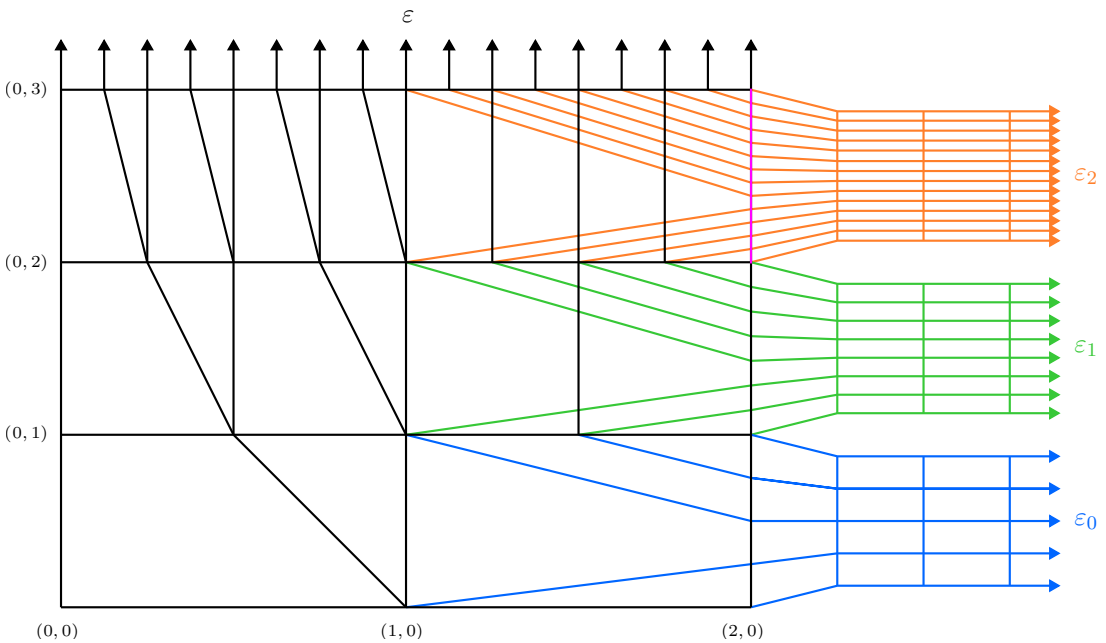

Figure 8.1: Depicted is the graph G from Construction 8.3.1. The subgraph induced by the orange edges is G_2 together with the extensions of its horizontal rays.

For our proofs of Examples 25 and 26 we construct a graph G in Construction 8.3.1 below, and then show that G is already as desired for both Examples 25 and 26. The graph in this construction is inspired by [CHM22, Example 7.4].[3]

[3]The example is only presented in the arXiv version of [CHM22].

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 $G_n := \mathbb{N}_{\geqslant 1} \square P_{2^{n+1}+2^n+1}$ 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 25 and 26. 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 [D17b, 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 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

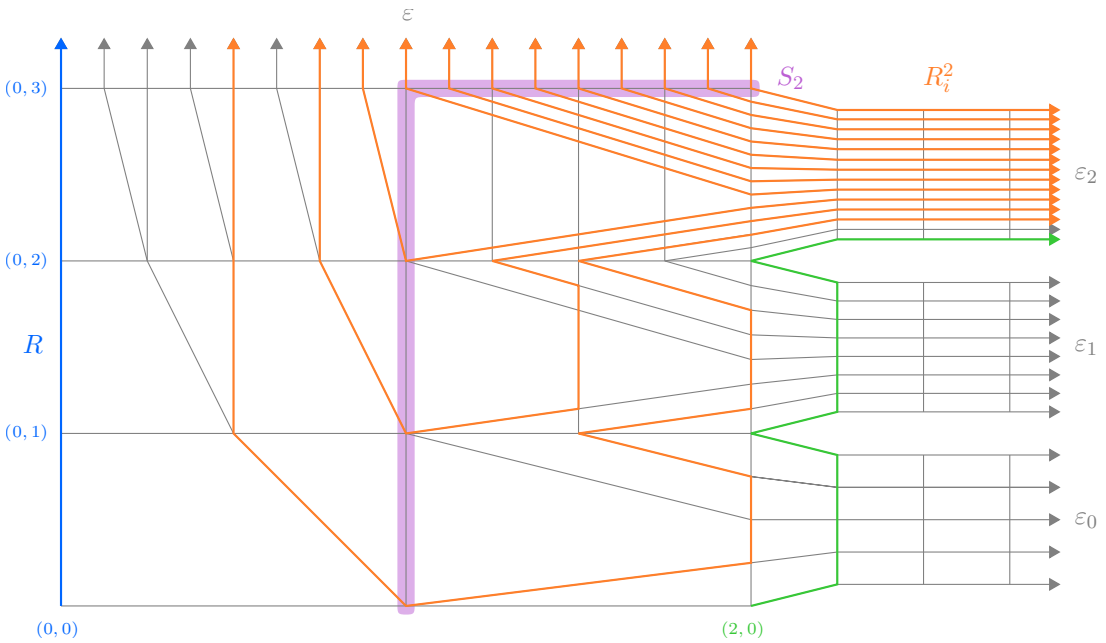

Figure 8.2: Depicted is the situation in the proof of Lemma 8.3.2 for $n = 2$.

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 $G \mathring{\uparrow} e$ 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 $G \mathring{\uparrow} e_n$ avoids the first column of G_n, the component D_n of $G \mathring{\uparrow} e_n$ in which ε_n lives is contained in G_n.

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

each row. Since there are $\deg(\varepsilon_n)$ rows, we have $|N_G(D_n)| \geqslant \deg(\varepsilon_n) = |V_{e_n}|$, and thus $N_G(D_n) = V_{e_n}$.

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

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

Recall that the adhesion set V_{e_n} is contained in G_n but avoids the first column of G_n, and thus $V_{e_n} \subseteq V(H_n)$. Since also $V_{e'_n} \cap V(H_n) = \emptyset$, there are at most $2^{n+1} + n + 2$ pairwise disjoint V_{e_n}–$V_{e'_n}$ paths in G, as witnessed by S_n. As $(T, \mathcal{V})$ is strongly linked by assumption, there is an edge f_n on the unique e_n–e'_n path in T such that $|V_{f_n}| \leqslant |S_n|$. Moreover, V_{f_n} separates V_{e_n} and $V_{e'_n}$, and thus also ε_n and ε. By the definition of G', there are in fact $|S_n|$ disjoint ε–ε_n double rays R^n_i in G (indicated in orange in Figure 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 26. 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 25, 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 25 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 [D17b, 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 25. 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 27 – negative answer to question (vi)

In this section we construct Example 27, which we restate here for convenience:

Example 27. *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 [C19, 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.

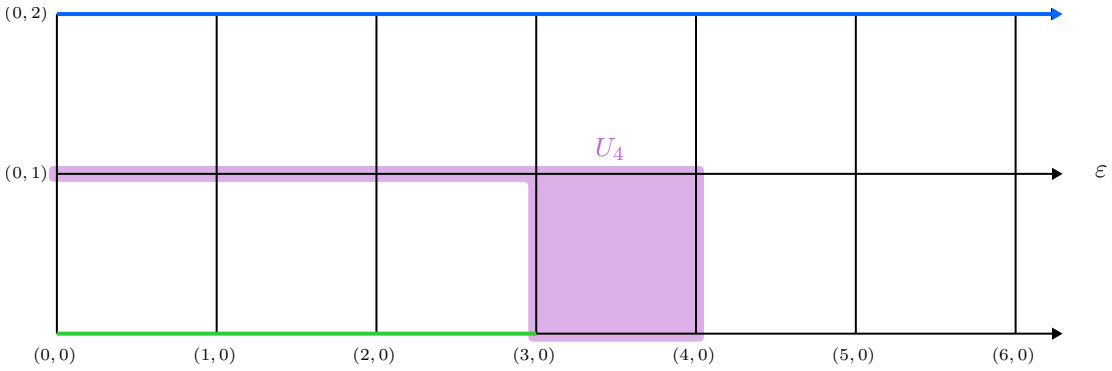

Figure 8.3: Depicted is the graph G' and the set U_4 in purple from Example 27. Indicated in blue is the ray $(0,2)(1,2)\dots$ that contains the vertex w and indicated in green is the path $(0,0)(1,0)(2,0)(3,0)$ that contains the vertex u in the case $m = 4$.

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

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

By the choice of f and m, we have $U_{m+1} \subseteq V(G{\uparrow}f)$, as U_{m+1} is complete. Hence, as V_f separates $G\mathring{\uparrow}f$ and $G\mathring{\downarrow}f$, it follows that $U_m \cap U_{m+1} \subseteq V_f \subseteq V_{t_1}$. Moreover, as $(0,2) \in V(G{\downarrow}e) \subseteq V(G{\downarrow}f)$ 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 [AJKP24, 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 [H75, 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 [H77, 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 $G{\uparrow}e'$, or is contained in a critical vertex set of G that is included in $G{\uparrow}e'$.

But (1) allows for more than (1'): If $e <_T e' \in T$ and $|V_e| > |V_{e'}|$, then V_e and $V_{e'}$ are allowed to intersect also in vertices that do not dominate an end and that are not contained in a critical vertex set. The following example shows that allowing this is in fact necessary.

[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.

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.

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

[6] We restrict to the set c' instead of $c_{00}(\mathbb{N})$ to ensure that G is connected.

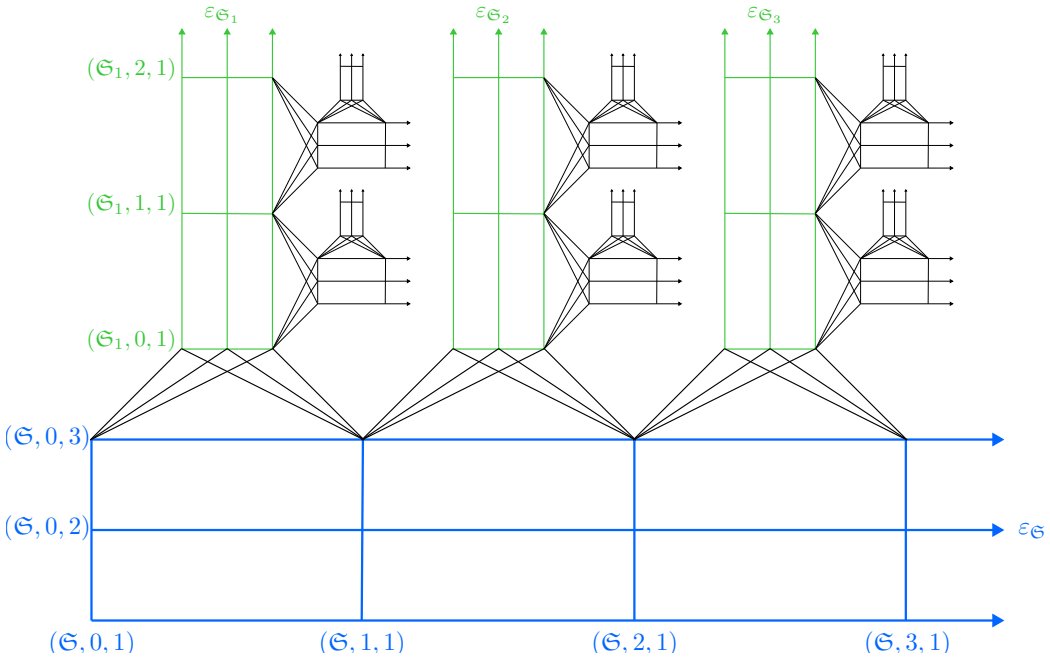

Figure 8.4: Depicted is the graph G from Example 8.5.2. The blue subgraph is induced by the vertices $(\mathfrak{S}, i, j)$ of G with $\mathfrak{S} = (0, 0, 0, \dots)$. The green subgraphs are induced by the vertices $(\mathfrak{S}_x, i, j)$ of G with $\mathfrak{S}_x = (x, 0, \dots)$ for $x \in \mathbb{N}_{\geqslant 1}$, respectively.

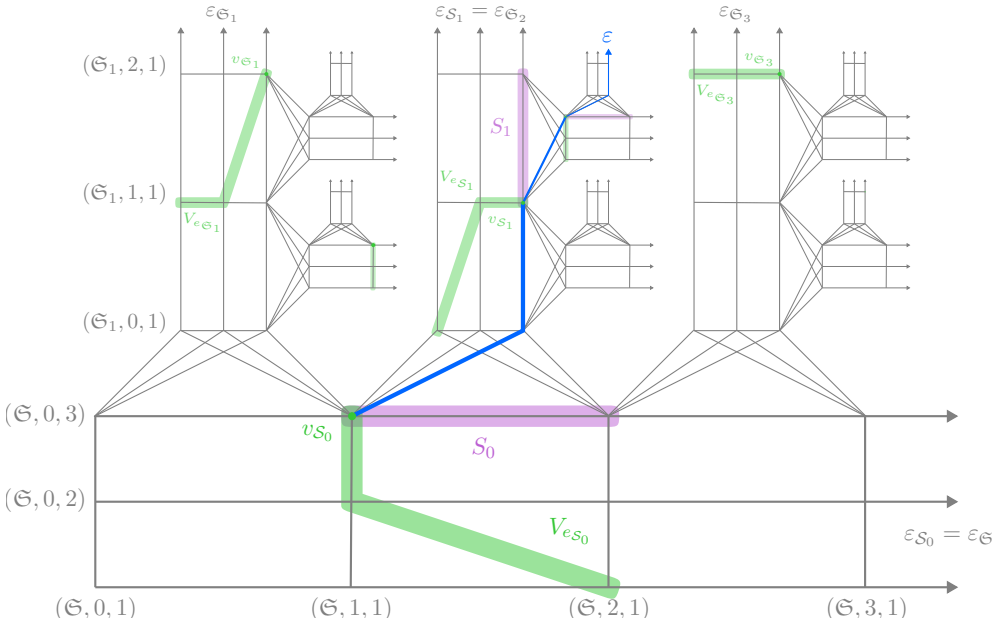

Figure 8.5: A ray in G that meets all the sets $V_{e_{\mathcal{S}_n}}$. The end to which it belongs has degree 2.

This example might explain why Halin never extended his result [H77, Satz 10] mentioned above to graphs with more than two ends: His precursor notion to a tree-decomposition, namely the quasi-trees and pseudo-trees discussed in [H91b], required upwards disjoint separators, and hence could not possibly capture the types of locally finite graphs described in Example 8.5.2.

CHAPTER 9

A grid theorem for strong immersions of walls

This chapter is based on [DJKW25] and joint work with Reinhard Diestel, Paul Knappe and Paul Wollan.

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.

9.1 Introduction

In their Graph Minors Project [RS83–04], 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; W15].

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 [W15] 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 [W15] 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 9.1.1 ([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 9.1.1, then there is also a qualitative converse of Theorem 9.1.1 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 9.1.1 with F restricted to walls, then Theorem 9.1.1 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 9.1). 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 9.1.1 hence yields an integer $\alpha > 0$ such 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.

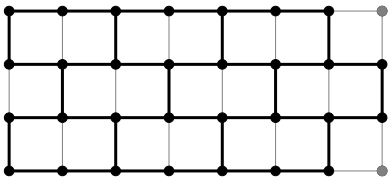

Figure 9.1: The wall W_4 of size 4 with the underlying 4×8 grid.

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 9.1.1 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 [W15], 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 9.2.

Our first main result, a version of Theorem 9.1.1 specifically tailored to walls, then reads as follows:

Theorem 28. *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 28 in that we cannot consider the torsos themselves instead of their 3-centres (see Example 9.2.6).

In contrast to Theorem 9.1.1, the notion of path-like torsos as in Theorem 28 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 28:

Theorem 29. *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 29 is self-contained and does not build on any previous results, the proof of Theorem 28 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; W15]. 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 9.2. Then we prove Theorem 29 in Section 9.3 and Theorem 28 in Section 9.4.

For basic graph-theoretic terms, we follow [D17b]. 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 [D17b] in that we define the *components* of a graph not as its maximal connected subgraphs, but as their vertex sets.

9.2 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 9.2.1. 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 9.2.2. *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 $V(G - X)$, 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,

[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.

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 9.2.3. *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 9.2.2, 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 9.2.4. *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 $\{v'_1,\dots,v'_{i-1}\}$ 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 28 (see Example 9.2.6). Formally, we achieve this with the notion of 3-centres which was introduced in [W15] and is well-defined by [W15, Lemma 9].

Definition 9.2.5. 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 9.2.4. In particular, it is weaker to assume in Theorem 29 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 9.2.2 and Example 9.2.3, 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 28 and consider the torsos themselves instead, as the following example demonstrates.

Example 9.2.6. *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 9.2.2, and hence H_s is not almost α-thin either. □

If we adapt Example 9.2.6 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.

9.3 Strongly immersed walls are obstructions

In this section, we prove Theorem 29 which we recall here:

Theorem 29. *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 29, we need Lemma 9.3.1, 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 9.3.1. *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 [M27], 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 $j = 1, \dots, \ell$ 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 $i = 1, \dots, \ell$ 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 29. 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 9.3.1, 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. □

9.4 Excluding strongly immersed walls

This section is devoted to the proof of Theorem 28, recalled here.

Theorem 28. *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 28, we first reduce the problem to 3-edge-connected graphs in Section 9.4.1. In Section 9.4.2 we collect some tools from previous research dealing with immersions and tree-cut decompositions. Then we finally prove Theorem 28 in Section 9.4.3.

9.4.1 Reduction to 3-edge-connected graphs

We begin by reducing the proof of Theorem 28 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 9.4.1. *If Theorem 28 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 28 holds for all graphs with minimum degree 3. We now show that Theorem 28 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 $t_ut_v\in T$ 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 9.4.2. *If Theorem 28 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 28 holds for all 3-edge-connected graphs. By Lemma 9.4.1, 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 $t \in T$.

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

[2] This construction follows the one in terms of edge-sums given in [W15, Lemma 5].

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 9.2.4. The analysis of the torso at t^B is symmetrical. □

9.4.2 Immersions and tree-cut decompositions

In this section we collect some tools and auxiliary results for the proof of Theorem 28.

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 28 we reduced to 3-edge-connected graphs in Section 9.4.1.

Lemma 9.4.3 ([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. [W15, Observation 1]).

Lemma 9.4.4. *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 9.4.5. *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 9.4.4. □

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 9.4.6 (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 9.4.7 ([W15, 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 9.4.8 ([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 28, 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 9.4.9 ([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.*

9.4.3 Proof of Theorem 28

Proof of Theorem 28. 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 9.4.2, we may assume that G is 3-edge-connected.

To construct the desired tree-cut decomposition $(T, \mathcal{X})$ of G, we first apply Lemma 9.4.7 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 9.4.7 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 9.4.6.*

Proof of the claim. 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 9.4.6 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 of the claim. 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 9.4.8 and Lemma 9.4.3, respectively. Since G' is 3-edge-connected and does not contain W_ℓ as strong immersion minor, Lemma 9.4.3 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 9.4.4. Thus, by Lemma 9.4.8, $A[C]$ 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 9.4.5, 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 9.4.7 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 $\alpha = \alpha(\ell)$ be defined as

$$\alpha(\ell) := \max\{k(\ell)(8\ell^2+1+6\ell^2 \cdot 16\ell^4), p(\ell) \cdot (6\ell^2)^2 + \binom{16\ell^4}{2} \cdot p(\ell) \cdot (6\ell^2)^2 + k(\ell) \cdot d(\ell)/2 + a(\ell) \cdot k(\ell)\}.$$

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 of the claim. 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 of the claim. 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 9.4.4, 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 9.4.9 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 9.4.5, 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 9.4.5, 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 28. □

Acknowledgements

We are very grateful to Nathan Bowler for the suggestion to bound the size of the neighbourhood of the vertices which we delete in the definition of almost α-thinness. This enabled us to obtain the actual equivalence given by Theorem 28 and Theorem 29 instead of an approximate one.

CHAPTER 10

The Lovász-Cherkassky theorem for locally finite graphs with ends

This chapter is based on [JJKKM23] and joint work with Attila Joó, Paul Knappe, Jan Kurkofka and Ruben Melcher.

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.

10.1 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 10.1.1. *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 10.1.2 (Lovász-Cherkassky Theorem [C77; L76])**.** *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*

$$\frac{1}{2}\sum_{t\in T}\lambda_G(t, T \setminus \{t\}).$$

In this note, we extend the theorem of Lovász and Cherkassky to all locally-finite infinite graphs and their ends. For infinite graphs and their ends, we follow and assume familiarity with the terminology in [D17b, §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 $\lambda_G(x, Y)$.

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. An end ω of a locally finite graph 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, Section 3] for a discussion of the parity of ends. A locally finite graph G is *inner-Eulerian* for a set $T \subseteq \hat{V}(G)$ if every vertex and end in $\hat{V}(G) \setminus T$ has even degree in G.

Our main result reads as follows:

Theorem 30. *Let G be any locally-finite connected 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 $\lambda_G(t, T \setminus \{t\})$.*

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 [D03, §3].)

We conclude the introduction with two examples that discuss why certain weakenings in the assumptions of Theorem 30 cannot be made.

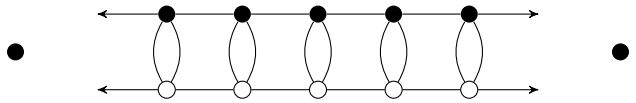

Figure 10.1: T consists of the black vertices and ends

Example 10.1.3. We claim that it is not possible to drop in Theorem 30 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 $|G|$. 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 10.1. 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.

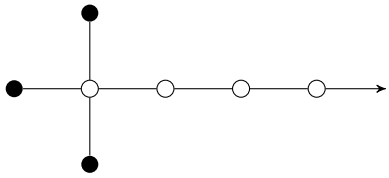

Figure 10.2: T consists of the black vertices

Example 10.1.4. Theorem 30 requires that all vertices and all ends in $\hat{V}(G) \setminus T$ have even degree in G. We claim that requiring even end degrees in addition to even vertex degrees is really necessary. Indeed, let us consider the tree G in Figure 10.2 and let T consist of its leaves. All the non-leaves of G have even degree, but the unique end of G has degree 1. Any set $\mathcal{A}$ of 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 $\lambda_G(t, T \setminus \{t\}) = 1$.

This note is organised as follows. In Section 10.2, we introduce the tools and terminology that we need. In Section 10.3, we prove Theorem 30.

10.2 Tools and terminology

In this section, we recall a number of results that we need in order to prove Theorem 30. Originally, all these results have been proved only for simple graphs. However, all of them

extend 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.

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 $v \in G$.

We will need the following generalisation of the Lovász-Cherkassky theorem to countably infinite graphs without ends:

Theorem 10.2.1 ([J23, 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 $\mathcal{P}$.*

In order to use Theorem 10.2.1, we will need the following lemma:

Lemma 10.2.2. *If G is a locally finite connected graph and $T \subseteq \hat{V}(G)$ is given such that every vertex and end in $\hat{V}(G) \setminus T$ has even degree, then 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.*

This lemma follows directly from the Infinite Handshaking Lemma, as follows.

Lemma 10.2.3 (Infinite Handshaking Lemma [BS07, Proposition 15])**.** *The number of odd vertices and ends in a locally finite graph is even or infinite.*

Proof of Lemma 10.2.2. Assume for a contradiction that there is a vertex set $X \subseteq V(G) \setminus T$ with $\overline{X} \cap T = \emptyset$, but such that $d_G(X)$ is an odd natural number. Then X and $V(G) \setminus X$ are non-empty. Consider the graph H that arises from G by contracting $G - X$ 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 10.2.3. □

Finally, we need the following lemma about the connectivity between an end of a graph and a finite set of vertices:

Lemma 10.2.4 ([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.*

10.3 Proof of Theorem 30

Proof of Theorem 30. 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 $(w_n \colon n < \kappa)$ 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–$(T \setminus \{\omega_n\})$

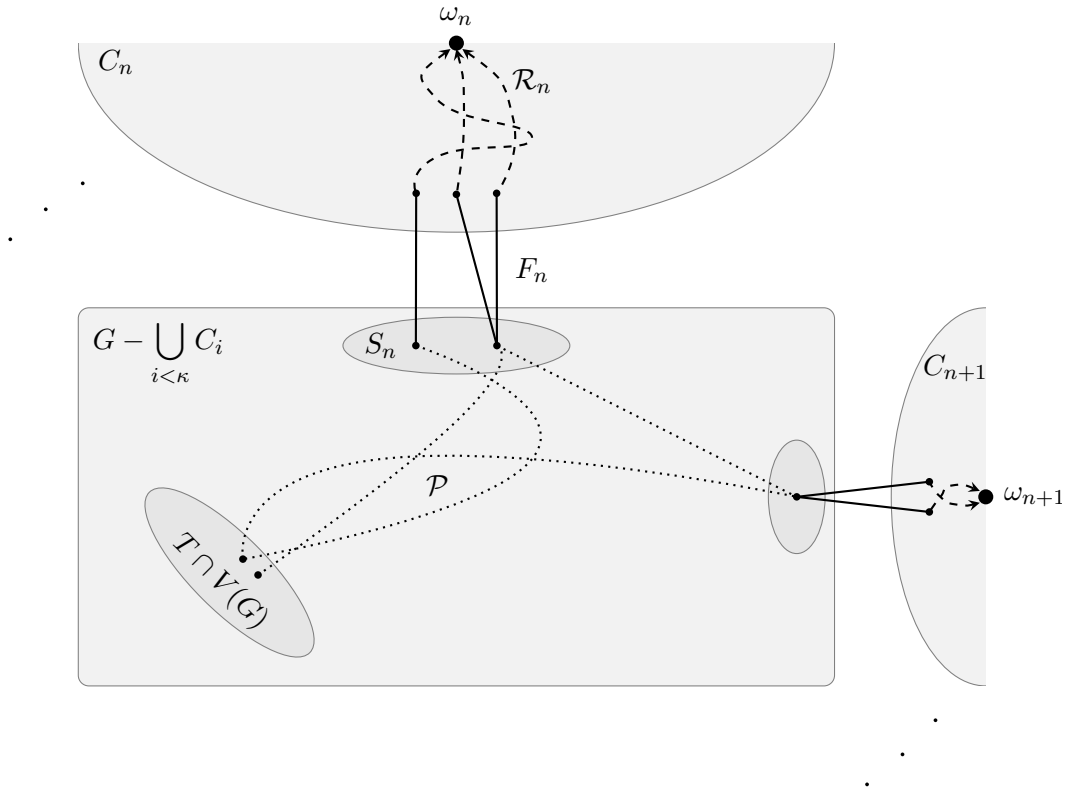

Figure 10.3: Situation in the proof of the main result

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 10.3 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 $|G_n| = |G| \,/\, \big\{\overline{C_i} \colon i < n\big\}$, 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.

By assumption, no vertex or end in $\hat{V}(G) \setminus T$ has odd degree in G. Hence, by Lemma 10.2.2, 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. It follows 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. And 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 10.2.1 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 10.2.4 to S_n and ω_n to find a set $\mathcal{R}_n$ of $|F_n|$ many edge-disjoint S_n–ω_n rays. Note that the rays in $\mathcal{R}_n$ 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 $\mathcal{R}_n$ 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 $\mathcal{R}_n$ 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 $\mathcal{R}_n$ and $\mathcal{R}_m$ 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 $\mathcal{R}_n$ 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. □

CHAPTER 11

Menger's theorem in bidirected graphs

This chapter is based on [BGGJR23] and joint work with Nathan Bowler, Ebrahim Ghorbani, Florian Gut and Florian Reich.

Bidirected graphs are a generalisation of directed graphs that arises in the study of undirected graphs with perfect matchings. Menger's famous theorem – the minimum size of a set separating two vertex sets X and Y is the same as the maximum number of disjoint paths connecting them – is generally not true in bidirected graphs. We introduce a sufficient condition for X and Y which yields a version of Menger's Theorem in bidirected graphs that in particular implies its directed counterpart.

11.1 Introduction

A fundamental result in the field of graph theory is due to Karl Menger [M27] which is nowadays just known as Menger's Theorem. It gives an important insight into the connectivity of two sets of vertices in directed graphs. In its common form, Menger's Theorem asserts that the minimum number of vertices separating two sets of vertices X and Y in a directed graph D is the same as the maximum number of disjoint X–Y paths in D, see e.g. [BG09, Theorem 7.3.1].

There are various proofs of Menger's Theorem (cf. [D17b, Section 3.3]). One by Böhme, Göring and Harant [BGH01] proves a slight inductive structural strengthening, sometimes called Perfect's Theorem [P68]: For a set $\mathcal{P}$ of k disjoint X–Y paths, there is either a set of k vertices separating X from Y, one vertex on each path in $\mathcal{P}$, or there is a set of $k+1$ disjoint X–Y paths where k of them use the same startvertices in X as the paths of $\mathcal{P}$ and k of them use the same endvertices in Y as the paths of $\mathcal{P}$.

Theorem 11.1.1 ([BGH01; P68])**.** *Let X and Y be sets of vertices of a directed graph D, and let $P_1, \dots, P_k$ be vertex-disjoint X–Y paths in D where P_i starts in $v_i \in X$ and ends in $w_i \in Y$ for $i \in [k]$. Then precisely one of the following is true:*

(1) There is a set S of k vertices of D such that $D - S$ contains no X–Y path.

(2) There are $k+1$ vertex-disjoint X–Y paths $P'_1, \dots, P'_{k+1}$ in D such that P'_i starts in v_i and some P'_j with $j \in [k+1]$ ends in w_i.

In this chapter we make Menger's Theorem available in the realm of bidirected graphs. Bidirected graphs were first introduced by Kotzig in a series of papers [K59a; K59b; K60].

These objects can be understood as a generalisation of undirected and directed graphs. They can be obtained from undirected graphs by assigning to each endpoint of every edge one of two signs, + or −. Directed graphs can then be envisioned as bidirected graphs in which each edge has sign + at one endpoint and − at the other.

Bidirected graphs have recently received increased attention by the graph theoretic community with both structural [AK06; W22] and algorithmic results [AFN96; BZ06]. Especially, Wiederrecht proposed bidirected graphs as another angle of attack to investigate the structure of undirected graphs with perfect matchings. This approach is inspired by a recent proof of Norin's Matching Grid Conjecture for bipartite graphs with perfect matchings, which was given by Hatzel, Rabinovich and Wiederrecht [HRW19]. Their proof applies the Directed Grid Theorem [KK15] using the structural relation of directed graphs and bipartite graphs with perfect matchings.

In the same structural sense, bidirected graphs are in a one-to-one correspondence with general undirected graphs with perfect matchings (for a detailed formal introduction, see [W22, Section 3.2]): Consider an undirected graph G with a perfect matching M. For each edge $e \in M$, we declare one of its endvertices as + and one as −; since M is a perfect matching, this gives each vertex of G a sign + or −. We then obtain a bidirected graph B from G and M by assigning an edge $e \notin M$ at its endpoint v the sign of v under the above declaration and then contracting all edges in M. Vice versa, a bidirected graph B gives rise to an undirected graph G with a perfect matching M by expanding each vertex $v \in B$ into an edge $e_v = v^+v^-$ and making all edges in B with sign + at v incident to v^+ and all edges with sign − at v incident to v^-. These two operations are inverse to each other, and they are structurally invariant under swapping the roles of + and − at a vertex/matching edge.

Under this correspondence, paths in the bidirected graph B correspond to M-alternating paths in the undirected graph G. Menger's Theorem in bidirected graphs therefore corresponds to Menger's Theorem for M-alternating paths in undirected graphs with a perfect matching M. In this way, Menger's Theorem in bidirected graphs becomes a fundamental tool to understand the connectivity structure of undirected graphs with perfect matchings in terms of their alternating paths, and as such it may form a next step towards a proof of Norine's Matching Grid Conjecture.

Unfortunately, Menger's Theorem as in Theorem 11.1.1 does not hold if we simply replace 'directed' by 'bidirected' as illustrated in Theorem 11.3.1 and Figure 11.1. The main reason for this is an intricate complication in the structure of bidirected graphs: 're-routing' between paths does not work in general, that is, when two paths meet in a vertex, the incident signs need not allow us to follow the initial segment of the first path up to the common vertex and then continue with the end segment of the second path. As a consequence, and unlike in (un)directed graphs, a walk between two fixed vertices need not contain a path between them. These properties of bidirected graphs in particular prevent a direct transfer of the usual proofs of Theorem 11.1.1 such as the one given in [BGH01].

To overcome this challenge, we introduce a sufficient condition on the set $\mathcal{X}$ of 'signed' startvertices of our paths, where a *signed vertex* is a pair of a vertex and a sign + or −: we require $\mathcal{X}$ to be 'clean', that is, we forbid the existence of nontrivial $\mathcal{X}$–$\mathcal{X}$ paths. This condition is satisfied in particular by any set of startvertices in a directed graph. Our main

result then reads as follows:

Theorem 31. *Let $\mathcal{X}$ and $\mathcal{Y}$ be sets of signed vertices of a bidirected graph B, and suppose that $\mathcal{X}$ is clean. Let $P_1, \dots, P_k$ be vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths in B where P_i starts in $v_i \in V(\mathcal{X})$ and ends in $w_i \in V(\mathcal{Y})$ for $i \in [k]$. Then precisely one of the following is true:*

(1) There is a set S of k vertices of B such that $B - S$ contains no $\mathcal{X}$–$\mathcal{Y}$ path.

(2) There are $k+1$ vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths $P'_1, \dots, P'_{k+1}$ in B such that P'_i starts in v_i and some P'_j with $j \in [k+1]$ ends in w_i.

The general strategy of our proof of Theorem 31 follows the idea of [BGH01], but needs several extensions in order to circumnavigate the above mentioned complication for paths in bidirected graphs. In particular, we introduce 'appendages' as our main tool to overcome the walk-path problem for clean sets of vertices. This allows us to first prove an edge-version of Menger's Theorem, Theorem 11.5.1, from which we then deduce the vertex-version, Theorem 31. Our proof of Theorem 31 can also be turned into a polynomial-time algorithm which finds a maximal set of disjoint paths.

Structure of the chapter: After introducing notation and general concepts around bidirected graphs in Section 11.2, we show in Section 11.3 that Menger's Theorem cannot be transferred directly to bidirected graphs. We then introduce 'appendages' in Section 11.4 and use it to first prove Menger's Theorem for edge-disjoint paths in Section 11.5. Finally, we deduce the vertex-version, Theorem 31, from the edge-disjoint one in Section 11.6. In Section 11.7 we describe how our proof of Theorem 31 can be turned into a polynomial-time algorithm.

11.2 Preliminaries

Throughout the chapter we assume $\mathbb{N} = \{0, 1, 2, \dots\}$. For numbers $k \leqslant \ell$ we write $[k, \ell] := \{n \in \mathbb{N} : k \leqslant n \leqslant \ell\}$, and we abbreviate $[k] := [1, k]$.

11.2.1 Basic Definitions for Bidirected Graphs

For the basic definitions concerning undirected graphs we refer the reader to [D17b] and for the notions around directed graphs we refer the reader to [BG09]. Undirected and directed graphs considered in this chapter are allowed to have parallel edges, but they do not have loops.

A *bidirected graph* $B = (G, \sigma)$ consists of a graph $G = (V, E)$, a corresponding set of *half-edges* defined as

$$\mathbf{E}(B) := \{(u, e) \colon e \in E \text{ and } u \in e\}$$

and a *signing* $\sigma : \mathbf{E} \to \{+, -\}$ assigning to each half-edge (u, e) its *sign* $\sigma(u, e) := \sigma((u, e))$; we say that e *has sign* $\sigma(u, e)$ *at* u. Then $V(B) := V$ is the *vertex set* of B and $E(B) := E$ is its *edge set*. We refer to the elements of $V(B)$ and $E(B)$ as the *vertices* and the *edges* of B, respectively.

For technical simplification, we do not allow bidirected graphs to have distinct edges e and f that have the same endvertices and the same signs at them. This simplification does not affect the main theorem, Theorem 31, which still holds without this assumption. Indeed, one can easily deduce the statement for bidirected graphs with parallel edges by applying Theorem 31 to the bidirected graph obtained by subdividing each edge and giving the two edges incident to a newly arising subdivision vertex distinct signs at it.

A *signed vertex* of B is a pair (v, α) of a vertex v of B and a sign $\alpha \in \{+, -\}$. Given a set $\mathcal{X}$ of signed vertices, we write $V(\mathcal{X}) := \{v \in V(B) \mid \exists\, \alpha \in \{+, -\} : (v, \alpha) \in \mathcal{X}\}$.

An *oriented edge* $\vec{e}$ of a bidirected graph B is formally defined as a triple (e, u, v) where e is an edge of B with incident vertices $u, v \in V(B)$; we call e its *underlying edge*, u its *startvertex* and v its *endvertex*, and think of $\vec{e}$ as orienting e from u to v. The edge e has precisely two *orientations*, one with startvertex u and endvertex v and the other one with startvertex v and endvertex u. We denote the two orientations of e as $\vec{e}$ and $\overleftarrow{e}$; there is no default orientation of e, but if we are given one of them as $\vec{e}$, say, then the other one is written as $\overleftarrow{e}$. Given an oriented edge $\vec{e}$ we conversely write e for its underlying edge. Given a set A of edges of B, we write $\overleftrightarrow{A}$ for the set of all orientations of edges in A.

In every figure involving bidirected graphs, we depict the signs of halfedges in a particular manner by drawing the signs onto the edge, which results in a bar perpendicular to the edge at incident vertices with sign + and none at incident vertices with sign −, see e.g. Figure 11.2.

11.2.2 Connectivity in Bidirected Graphs

Let $B = (G, \sigma)$ be a bidirected graph. A sequence

$$W = v_0 \vec{e}_1 v_1 \vec{e}_2 v_2 \dots v_{\ell-1} \vec{e}_\ell v_\ell$$

of vertices $v_i \in V(B)$ and oriented edges $\vec{e}_j \in \vec{E}(B)$ is a *walk* W of *length* ℓ in B if the oriented edge $\vec{e}_j$ has startvertex v_{j-1} and endvertex v_j for $j \in [\ell]$ and we have $\sigma(v_i, e_i) \neq \sigma(v_i, e_{i+1})$ for $i \in [\ell - 1]$. A *subwalk* of W is a walk of the form $v_i \vec{e}_i \dots \vec{e}_{j-1} v_j$ for some $i \leqslant j$. The *inverse of* W is $W^- := v_\ell \overleftarrow{e}_\ell v_{\ell-1} \overleftarrow{e}_{\ell-1} v_{\ell-2} \dots v_1 \overleftarrow{e}_1 v_0$. A walk is *trivial* if it has length 0 and *nontrivial* otherwise.

The walk W *starts* in its *startvertex* v_0 and *ends* in its *endvertex* v_ℓ. If W is nontrivial, we say that it starts with sign $\sigma(e_1, v_0)$ in its *signed startvertex* $(v_0, \sigma(e_1, v_0))$. Likewise we say that it ends with sign $\sigma(e_\ell, v_\ell)$ in its *signed endvertex* $(v_\ell, \sigma(e_\ell, v_\ell))$. All other vertices of W are *internal vertices* of W. The set of all vertices in W is denoted as $V(W) := \{v_0, \dots, v_\ell\}$. The nontrivial walk W *arrives in* v_i with sign $\sigma(e_i, v_i)$ for $i \in [1, \ell]$ and *leaves* v_i with sign $\sigma(e_{i+1}, v_i)$ for $i \in [0, \ell - 1]$. Analogously, we say that a nontrivial walk W *starts* in $\vec{e}_1$ and *ends* in $\vec{e}_\ell$.

The set of all oriented edges in a walk W is denoted by $\vec{E}(W) := \{\vec{e}_1, \dots, \vec{e}_\ell\}$, the set of edges underlying $\vec{E}(W)$ is $E(W) := \{e_1, \dots, e_\ell\}$, and we write $\overleftrightarrow{E}(W) := \overleftrightarrow{E(W)}$. If $E(W) \subseteq A$ for some set A of edges of B, then W is a walk *in* A.

A nontrivial walk W is an *x–y walk* in B if W starts in x and ends in y, where x and y can be oriented edges or vertices or signed vertices of B. Given two sets $\mathcal{X}$ and $\mathcal{Y}$ of signed vertices of B we define an *$\mathcal{X}$–$\mathcal{Y}$ walk* as follows: a trivial walk $W = v$ is an $\mathcal{X}$–$\mathcal{Y}$ walk

if there is $\alpha \in \{+,-\}$ such that $(v,\alpha) \in \mathcal{X}$ and $(v,-\alpha) \in \mathcal{Y}$. A nontrivial walk W is an $\mathcal{X}$–$\mathcal{Y}$ walk if it is an x–y walk for some $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, no internal vertex of W is in $V(\mathcal{X}) \cup V(\mathcal{Y})$ and neither the start- nor the endvertex of W forms a trivial $\mathcal{X}$–$\mathcal{Y}$ walk. Note that no proper subwalk of an $\mathcal{X}$–$\mathcal{Y}$ walk is again an $\mathcal{X}$–$\mathcal{Y}$ walk.

Two walks W and W' in B are *vertex-disjoint* if $V(W) \cap V(W') = \emptyset$. Similarly, W and W' are *edge-disjoint* if $E(W) \cap E(W') = \emptyset$. A set of walks is *vertex-disjoint* (respectively *edge-disjoint*) if its elements are pairwise vertex-disjoint (respectively edge-disjoint).

A walk W in B is a *trail* if all edges underlying $\vec{E}(W)$ are distinct. If not only all edges, but also all vertices of W are distinct, then W is a *path*. For both trails and paths we transfer all notions and notation defined for walks (such as e.g. subtrail or v–w path).

A key difference between bidirected and both undirected and directed graphs which complicates understanding their connectivity properties considerably lies in the relation of walks, trails and paths: unlike for (un)directed graphs, the existence of a v–w walk between two vertices v and w in a bidirected graph B does not imply that there exists a v–w trail in B, and similarly a v–w trail does not guarantee the existence of a v–w path (see [W22, Figure 9.1] for examples).

For a given trail $Q = v_0\vec{e}_1 v_1 \vec{e}_2 v_2 \dots v_{\ell-1}\vec{e}_\ell v_\ell$ and $0 \leqslant i \leqslant j \leqslant \ell$, we write $\vec{e}_i Q := v_{i-1}\vec{e}_i \dots v_\ell$, $Q\vec{e}_i := v_0 \dots \vec{e}_i v_i$ and $\vec{e}_i Q \vec{e}_j := v_{i-1}\vec{e}_i \dots \vec{e}_j v_j$ for the appropriate subtrails of Q. Analogously, given a path $P = v_0\vec{e}_1 v_1 \vec{e}_2 v_2 \dots v_{\ell-1}\vec{e}_\ell v_\ell$ and $0 \leqslant i \leqslant j \leqslant \ell$, we write $v_i P := v_i \dots v_\ell$, $P v_i := v_0 \dots v_i$ and $v_i P v_j := v_i \dots v_j$ for the appropriate subpaths of P, and let $\vec{e}_i P v_j := v_{i-1} P v_j$ as well as $v_i P \vec{e}_j := v_i P v_j$. In particular, these two notions combine for paths to the respective notations $\vec{e}_i P v_j$ and $v_i P \vec{e}_j$.

We use an analogous notion for the concatenation of trails at edges: for example if the union $Q_1\vec{e} \cup \vec{e}Q_2\vec{f} \cup \vec{f}Q_3$ of three trails is again a trail, we denote it as $Q_1\vec{e}Q_2\vec{f}Q_3$. We write similarly $P_1 x P_2 y P_3$ for paths where x, y may be both vertices or oriented edges.

The complicated connectivity structure of bidirected graphs also manifests in various different notions of strong connectivity (see [W22, Section 9.2] for an overview). We discuss two of them, namely 'strongly connected' and 'circularly connected', in the remainder of this section. Let us first define 'strongly connected':

Definition 11.2.1. We say a bidirected graph B is *strongly connected* if for any two vertices v and w of B there are signs α and β such that B contains both a (v,α)–(w,β) path and a $(v,-\alpha)$–$(w,-\beta)$ path.

It turns out that an apparently weaker condition is equivalent to the one in this definition:

Lemma 11.2.2. *Let B be a bidirected graph such that for any two vertices v and w of B there are a $(v,+)$–w path and a $(v,-)$–w path in B. Then B is strongly connected.*

Proof. Let v and w be vertices of B. We must find paths joining them as in the definition of strong connectivity. By assumption there are signs β_1 and β_2 such that B contains a $(v,+)$–(w,β_1) path and a $(v,-)$–(w,β_2) path. Similarly, applying the assumption to w and v then reversing the paths shows that there are signs α_1 and α_2 such that B contains a (v,α_1)–$(w,+)$ path and a (v,α_2)–$(w,-)$ path.

If $\beta_2 = -\beta_1$ then we are done, so we may suppose $\beta_1 = \beta_2$. Similarly if $\alpha_2 = -\alpha_1$ then we are done, so we may suppose $\alpha_1 = \alpha_2$. Setting $\alpha := -\alpha_1$ and $\beta := \beta_1$ we are done. □

We now turn to 'circularly connected'. For this, a *cycle* in a bidirected graph $B = (G, \sigma)$ is a trail $C = v_0\vec{e}_1 v_1\vec{e}_2 v_2 \dots v_{\ell-1}\vec{e}_\ell v_\ell$ in B whose vertices are all distinct except $v_0 = v_\ell$ where we also have $\sigma(v_0, e_1) \neq \sigma(v_\ell, e_\ell)$.

Definition 11.2.3. Let $B = (G, \sigma)$ be a bidirected graph and consider the undirected graph $H := (V(G), F)$ where F is the set of edges of G that lie on some cycle of B. We refer to the connected components of H as the *circular components* of B. If H is connected we call B *circularly connected*.

We now aim to show that 'strongly connected' and 'circularly connected' in fact describe the same type of connectivity. This result was already known (see for example [W22, Section 9.2]), but we include a proof here for the convenience of the reader.

Theorem 11.2.4. *A bidirected graph $B = (G, \sigma)$ is strongly connected if and only if it is circularly connected.*

For the proof of Theorem 11.2.4, we make use of a lemma that describes which types of connectivity are forced by the interaction of paths between two vertices that differ in the signs at their endpoints:

Lemma 11.2.5. *Let v and w be vertices of a bidirected graph B. Let α and β be signs such that B contains both a (v, α)–(w, β) path P and a $(v, -\alpha)$–$(w, -\beta)$ path Q. Then at least one of the following statements holds:*

- *v and w lie in the same circular component of B.*
- *There is a (v, α)–$(w, -\beta)$ path in B.*

Proof. The proof proceeds by induction on the sum of the lengths of the paths P and Q. If this sum is 0, then $v = w$ and therefore the first statement holds. Otherwise, let x be the first vertex of P other than v which also lies on Q. Suppose that xP is an (x, γ)–(w, β) path and that xQ is an (x, δ)–$(w, -\beta)$ path.

Case 1: $\gamma = \delta$ In this case, PxQ is a (v, α)–$(w, -\beta)$ path in B.

Case 2: $\gamma = -\delta$ In this case we can apply the induction hypothesis to x, w, γ, β and the paths xP and xQ in the bidirected graph B' given by the union of xP and xQ. There are two possibilities.

Case 2.1: x **and** w **lie in the same circular component of** B As the path Px is a (v, α)–$(x, -\gamma)$ path, Qx is a $(v, -\alpha)$–(x, γ) path, and these paths are disjoint except at their endvertices, their union is a cycle. This cycle witnesses that v is in the same circular component as x, hence also as w.

Case 2.2: There is an (x, γ)–$(w, -\beta)$ **path** R **in** B' In this case, PxR is a (v, α)–$(w, -\beta)$ path in B.

□

Proof of Theorem 11.2.4. Suppose first that B is circularly connected. By Lemma 11.2.2 it is enough to show that for any vertices v and w of B there are both a $(v,+)$–w path and a $(v,-)$–w path. Let $v \in V(B)$ and $\alpha \in \{+,-\}$ be arbitrary. We set $X \subseteq V(B)$ to be the set of all vertices x for which there is a (v,α)–x path. Note that the set X is non-empty since it contains v. We show that X is the whole vertex set of B. Suppose not, then by circular connectivity there is some edge e of B contained in some cycle C of B which joins some $x \in X$ and some $y \notin X$.

There is a (v,α)–x path P by definition of X. Let z be the first vertex of P on C and let $\beta \in \{+,-\}$ such that the (possibly trivial) path Pz ends in (z,β). Then there is a $(z,-\beta)$–y path Q contained in C. Thus, PzQ is a (v,α)–y path, contradicting $y \notin X$. Since this is true for any $v \in V(B)$ and any $\alpha \in \{+,-\}$, this completes this direction of the proof.

Now suppose instead that B is strongly connected, and suppose for a contradiction that it is not circularly connected. Then it must contain an edge e which joins two vertices x and y in different circular components. Let $\sigma(x,e) = -\alpha$ and $\sigma(y,e) = \beta$. Suppose there is no (x,α)–$(y,-\beta)$ path. Since B is strongly connected, there must be an (x,α)–(y,β) path and an $(x,-\alpha)$–$(y,-\beta)$ path. Applying Lemma 11.2.5, we yield an (x,α)–$(y,-\beta)$ path in B. That path together with e is a cycle in B, contradicting our assumption that x and y lie in different circular components. □

11.3 Counterexample to a general Menger-type theorem

In this section we show that Menger's Theorem, Theorem 11.1.1, does not hold true in bidirected graphs if we transfer the statement verbatim. Even if we allow the separating set S to have any fixed size k for some $k \in \mathbb{N}$, the statement is false:

Theorem 11.3.1. *For each number $k \in \mathbb{N}$, there exists a bidirected graph B and disjoint sets $\mathcal{X}$ and $\mathcal{Y}$ of signed vertices of B, each of size at least k, such that there are no two disjoint $\mathcal{X}$–$\mathcal{Y}$ paths and for each subset $S \subseteq V(B)$ of size k there exists an $\mathcal{X}$–$\mathcal{Y}$ path in $B - S$.*

For the proof of Theorem 11.3.1 we will rely on the following well-known topological lemma, which follows directly from the fact that a complete graph on five vertices is not planar:

Lemma 11.3.2. *Let x_1, x_2, y_1 and y_2 be points appearing in the clockwise order on the boundary of the closed disc. Then any arc from x_1 to y_1 meets any arc from x_2 to y_2.* □

Proof of Theorem 11.3.1. The vertices of B will be given by vertices $x_{i,j}$ indexed by pairs (i,j) of natural numbers with $i+j \in [2k+1]$. We add edges of three kinds. For any natural numbers i and j with $1 \leqslant i$ and $i+j \in [2k]$ we add an edge that is incident to $x_{i,j}$ with sign $-$ and incident to $x_{i,j+1}$ with sign $+$. Similarly for any natural numbers i and j with $1 \leqslant j$ and $i+j \in [2k]$ we add an edge that is incident to $x_{i,j}$ with sign $-$ and incident to $x_{i+1,j}$ with sign $+$. Finally, for any $i \in [0,2k]$ we add an edge incident to $x_{i,2k+1-i}$ and $x_{i+1,2k-i}$, with sign $-$ at both vertices. We call the edges of these three kinds *vertical*, *horizontal* and *diagonal edges* respectively – see Figure 11.1 for an illustration where $k = 2$. We set $\mathcal{X} := \{x_{0,j} \colon j \in [2k+1]\}$ and $\mathcal{Y} := \{x_{i,0} \colon i \in [2k+1]\}$.

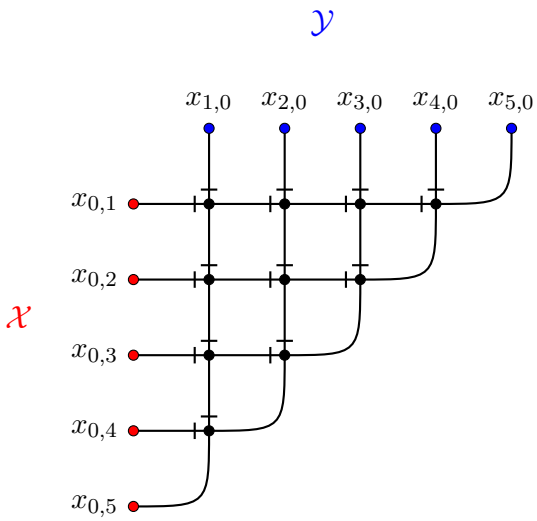

Figure 11.1: A bidirected graph containing neither two vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths nor a vertex set W of size $\leqslant 2$ such that in $B - W$ there are no $\mathcal{X}$–$\mathcal{Y}$ paths.

We show that for any set $S \subseteq V(B)$ of size k there is an $\mathcal{X}$–$\mathcal{Y}$ path in $B - S$. First we construct a sequence of $\mathcal{X}$–$\mathcal{Y}$ paths $P_1, P_2, \ldots, P_{2k+1}$ by taking the path P_i to be the concatenation of the path from $x_{0,i}$ to $x_{2k+1-i,i}$ consisting of horizontal edges, followed by the diagonal edge from $x_{2k+1-i,i}$ to $x_{2k+2-i,i-1}$ and then the path consisting of vertical edges to $x_{2k+2-i,0}$. Note that no vertex appears in more than two of these paths, so the number of these paths meeting S is at most $2k$. Since there are $2k+1$ of these paths, one of them avoids S.

Now suppose for a contradiction that there are two disjoint $\mathcal{X}$–$\mathcal{Y}$ paths P and Q. Let the initial vertices of P and Q be x_P and x_Q and let their final vertices be y_P and y_Q. Since deleting all the diagonal edges from B leaves a directed graph in which no directed edge points into $\mathcal{X}$ or $\mathcal{Y}$, both P and Q must contain diagonal edges. Let z_P and z_Q be endvertices of those diagonal edges on P and Q respectively which are distinct from x_P, x_Q, y_P and y_Q. We may embed B in the disc in such a way that on the boundary we have, in clockwise order, first the elements of $\mathcal{X}$, then those of $\mathcal{Y}$, then all other $x_{i,j}$ with $i + j = 2k + 1$. Without loss of generality we may assume that in the clockwise order after $\mathcal{Y}$ on the boundary of the disc we first reach z_Q then z_P. Then disjointness of the arcs induced by the paths $z_P P y_P$ and $x_Q Q z_Q$ contradicts Lemma 11.3.2, completing the proof. □

We will show in Section 11.6 that the vertex-version of Menger's Theorem in the context of bidirected graphs follows from the edge-version, meaning that the counterexample above also implies the existence of a counterexample to the edge-version of Menger's Theorem in our context. It is also not difficult to construct counterexamples to the vertex-version directly with a slight modification of the construction above, as illustrated in Figure 11.2.

Theorem 11.3.3. *For each number $k \in \mathbb{N}$, there exists a bidirected graph B and distinct vertices x and y, such that there are no two edge-disjoint x–y paths and for each subset $S \subseteq E(B)$ of size k there exists an x–y path in $B - S$.* □

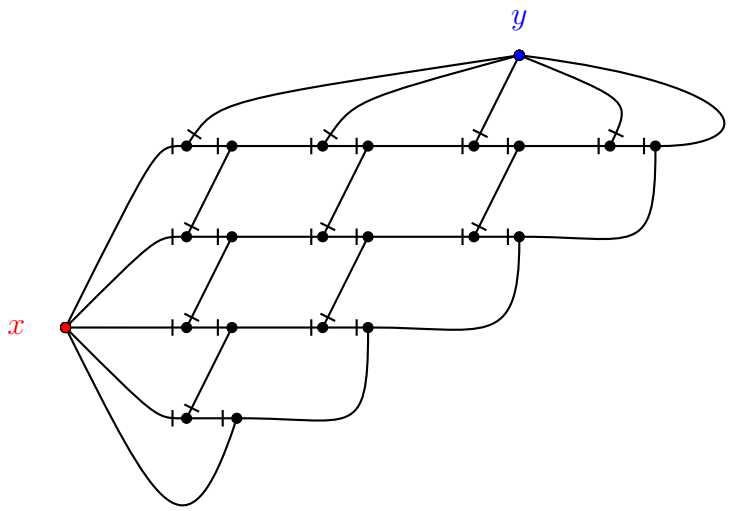

Figure 11.2: A bidirected graph containing neither two edge-disjoint x–y paths nor an edge set W of size $\leqslant 2$ such that in $B - W$ there is no x–y path.

11.4 Appendages of paths

In this section, we introduce the main tool for the proof of the edge-version of Menger's Theorem, Theorem 11.5.1. Given an arbitrary path P of a bidirected graph B starting in a vertex x, we define a set $A(P, x)$ of edges that includes much of the structure of B in the vicinity of P:

Definition 11.4.1. An edge set $A \subseteq E(B)$ is (P, x)-*admissible* if for any $\vec{e} \in \overleftrightarrow{A}$ there is an $\vec{e}$–x trail in $A \cup E(P)$. The *appendage* $A(P, x)$ of P and x is the union of all (P, x)-admissible sets.

Remark 11.4.2. Any union of (P, x)-admissible sets is again (P, x)-admissible. Thus, $A(P, x)$ is the maximal (P, x)-admissible set.

In the vertex-version of Menger's Theorem, Theorem 31, we restrict to bidirected graphs B and sets $\mathcal{X}$ of signed vertices such that B contains no nontrivial path that starts and ends in $\mathcal{X}$. However, for an edge-version of Menger's Theorem, forbidding such paths is not sufficient: joining the counterexample of Theorem 11.3.3 with a bidirected graph akin to the one depicted in Figure 11.3 by identifying the graphs at the respective vertices called x gives rise to a bidirected graph with no x'–x' trails whose internal vertices are all distinct[1]. But there are neither two disjoint x'–y paths nor is there a small set S of edges such that $B - S$ contains no x'–y path.
For the edge-version of Menger's Theorem, we thus forbid not only nontrivial x–x trails whose internal vertices are all distinct, but even nontrivial x–x trails in general:

Definition 11.4.3. A vertex x of a bidirected graph B is *edge-clean* if there exists no nontrivial x–x trail in B.

[1]We use this term rather than 'x'–x' path', as the definition of a path requires all vertices to be distinct.

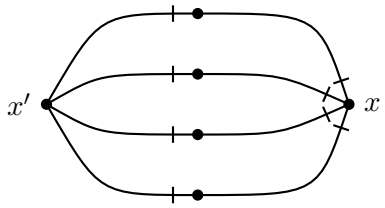

Figure 11.3: A bidirected graph with no x'–x' trail whose internal vertices are all distinct and the property that for any edge set W of size at most 3 there exists an x'-x path in $B - W$.

From now on we assume x to be edge-clean. Our aim is to show three essential properties of the appendage $A(P, x)$ that makes it a key tool in our proof of Theorem 11.5.1, the edge-version of Menger's Theorem in bidirected graphs. First, $A(P, x)$ is edge-disjoint from all paths starting in x that are edge-disjoint from P. Second, for any $v \in V(A(P, v)) \cup V(P)$ and any $\alpha \in \{+, -\}$ if there is an x–(v, α) path in B then there also is one in $A(P, x) \cup E(P)$ that coincides with P in the first edge. Third, any x–y walk contains an x–y path if it can be partitioned into a path in $A(P, x) \cup E(P)$ and a path in $E(B) \setminus (A(P, x) \cup E(P))$.

In the proof of Theorem 11.5.1, we replace the set $A(P, x) \cup E(P)$ by some auxiliary edges for a fixed path P. In this construction we do not want to remove paths starting in x that are edge-disjoint to P. Therefore it is essential that all paths starting in x which are disjoint to $E(P)$ are also disjoint to $A(P, x)$:

Lemma 11.4.4. *Let x be edge-clean and let Q be a path in B starting in x that is edge-disjoint from P. Then $E(Q) \cap A(P, x) = \emptyset$.*

Proof. Suppose not for a contradiction, and let $\vec{e}$ be the first edge of Q in $A(P, x)$. Since $A(P, x)$ is (P, x)-admissible, there is an $\vec{e}$–x-trail R in $A(P, x) \cup E(P)$. But then $Q\vec{e}R$ is a nontrivial x–x trail, contradicting the edge-cleanness of x. □

Corollary 11.4.5. *Let x be edge-clean. Then $A(P, x) \cup E(P)$ contains exactly one edge incident to x.*

Proof. Suppose not and let $e \in A(P, x) \setminus E(P)$ be incident to x. Then the path consisting of the single edge e contradicts Lemma 11.4.4. □

Now we turn our attention to the second essential property of appendages, which provides the existence of specific paths in $A(P, x) \cup E(P)$. We use this property in the proof of the edge-version of Menger's Theorem to ensure that certain paths are contained in $A(P, x) \cup E(P)$ by redirecting them if necessary.

Lemma 11.4.6. *Let x be edge-clean and let (v, α) be a signed vertex with $v \in V(A(P, x)) \cup V(P)$. Suppose that there is an x–(v, α) path in B. Then there is such a path in $A(P, x) \cup E(P)$.*

Proof. Let Q be an x–(v, α) path in B chosen to minimise its set of edges outside $A(P, x) \cup E(P)$. We show that Q is contained in $A(P, x) \cup E(P)$. Suppose not for a contradiction, and let $\vec{e}$ be the first edge of Q outside $A(P, x) \cup E(P)$.

Case 1: Some vertex of Q after $\vec{e}$ lies on P In this case, let w be the first such vertex along P. Note that $w \neq x$ since Q is a path. Let p be the edge preceding w on P and let q be the edge preceding w on Q.

Case 1.1: $\sigma(p,w) = \sigma(q,w)$ The path PwQ is an x–(v,α) path using fewer edges outside $A(P,x) \cup E(P)$ than Q, since Pw is disjoint from wQ per choice of w and $e \notin E(PwQ)$. This contradicts the minimality of Q.

Case 1.2: $\sigma(p,w) = -\sigma(q,w)$ We will show that the set $A(P,x) \cup E(Qw)$ is (P,x)-admissible. That is, for any oriented edge $\vec{f}$ with $f \in A(P,x) \cup E(Qw)$ there is an $\vec{f}$–x trail in $A(P,x) \cup E(Qw) \cup E(P)$. If $f \in A(P,x)$ this is clear, so suppose not. Thus, f is an edge of $\vec{e}Qw$. If $\vec{f} \in \vec{E}(Q)$ then $\vec{f}QwP^-$ is a suitable trail, by choice of w. If $\overleftarrow{f} \in \vec{E}(Q)$ then $\vec{f}\,Q^-$ is a suitable trail. Thus, $A(P,x) \cup E(Qw)$ is (P,x)-admissible, and so it is a subset of $A(P,x)$. This contradicts the assumption that e is not contained in $A(P,x)$.

Case 2: No vertex of Q after $\vec{e}$ lies on P In this case, let w be the first vertex along Q after $\vec{e}$ in $V(A(P,x)) \cup V(P)$. There is such a vertex since v is a candidate. Let q be the edge preceding w on Q. Our first aim is to show that there is $\vec{a} \in \overleftrightarrow{A(P,x)}$ with startvertex w and such that $\sigma(a,w) = -\sigma(q,w)$. Indeed, as $w \in V(A(P,v))$ there is $\vec{b} \in \overleftrightarrow{A(P,x)}$ so that its endvertex is w. If $\sigma(b,w) = -\sigma(q,w)$ then we can just set $\vec{a} := \overleftarrow{b}$. Otherwise, we can take $\vec{a}$ as the second edge of any $\vec{b}$–x trail in $A(P,x) \cup E(P)$. Let R be any $\vec{a}$–x trail in $A(P,x) \cup E(P)$.

We will now show that $A(P,x) \cup E(Qw)$ is (P,x)-admissible, that is, for any oriented edge $\vec{f}$ with $f \in A(P,x) \cup E(Qw)$, there is an $\vec{f}$–x trail in $A(P,x) \cup E(Qw) \cup E(P)$. If $f \in A(P,x)$ this is clear, so suppose not. Thus, f is an edge of $\vec{e}Qw$. If $\vec{f} \in \vec{E}(Q)$, then we find a suitable trail by following Q from $\vec{f}$ to w and switching onto $\vec{a}R$. If $\overleftarrow{f} \in \vec{E}(Q)$, then $\vec{f}\,Q^-$ is a suitable trail. Thus, $A(P,x) \cup E(Qw)$ is (P,x)-admissible, and so it is a subset of $A(P,x)$. This contradicts our assumption that e is not contained in $A(P,x)$.

This completes the proof. □

In the proof of Theorem 11.5.1, we will find some of the desired x–y paths in certain x–y walks. As mentioned in the introduction, an x–y walk in a bidirected graph does not necessarily contain an x–y path. We here provide a sufficient condition for the existence of an x–y path in an x–y walk, which uses the above notion of appendage for a given path P starting in x:

Lemma 11.4.7. *Let y and z be vertices of B such that x, y, z are pairwise distinct, and let $\alpha \in \{+,-\}$. Further, let Q be an x–(z,α) path that is internally disjoint to y such that all edges of Q except possibly the last edge are contained in $E(P) \cup A(P,x)$. Moreover, let R be a $(z,-\alpha)$–y path that is internally disjoint to x and avoids $E(P) \cup A(P,x)$. Then there is an x–y path S in $E(Q) \cup E(R)$ such that the first edge of S and the first edge of Q coincide, and the last edge of S and the last edge of R coincide.*

Proof. Let a be the first vertex of Q in $V(R)$. If $a = z$, then Q and R intersect only in $a = z$ and therefore QzR is the desired path. Note that $a \neq x$, since R is internally disjoint

to x. Thus, we can assume that a is an internal vertex of Q. Let e be the edge of Q preceding a and let f be the edge of R succeeding a. We show that $\sigma(a,e) = -\sigma(a,f)$ holds. Then $S := QaR$ forms an x–y path in $E(Q) \cup E(R)$. In particular, the path Qa is nontrivial by the choice of a and aR is nontrivial since y is not an internal vertex of Q. Thus, the first edge and the last edge of S are as desired.

Suppose for a contradiction that $\sigma(a,e) = \sigma(a,f)$ holds. Then aR^-zQ^-x and $aQzRaQ^-x$ are trails. We show that $E(Ra) \subseteq A(P,x)$, which contradicts the assumption that R avoids $A(P,x)$. More precisely, we prove that $A(P,x) \cup E(aQ) \cup E(Ra)$ is (P,x)-admissible: Let $\vec{e} \in \overleftarrow{E(aQ)} \cup \overleftarrow{E(Ra)}$. Then $\vec{e} \in \vec{E}(aR^-zQ^-x)$ or $\vec{e} \in \vec{E}(aQzRaQ^-x)$ holds. Note that both trails are contained in $A(P,x) \cup E(aQ) \cup E(Ra)$, as $E(Qa) \subseteq A(P,x)$ holds since $a \neq z$. This completes the proof. □

11.5 Menger's theorem for edge-disjoint paths

In this section we prove an edge-version of Menger's Theorem in bidirected graphs, from which we will then deduce our main result, the vertex-version given by Theorem 31, in the subsequent Section 11.6:

Theorem 11.5.1. *Let x and y be distinct vertices of a bidirected graph B, and suppose that x is edge-clean. Let $P_1, \dots, P_k$ be edge-disjoint x–y paths in B where $\vec{e}_i$ is the first edge of P_i and $\vec{f}_i$ is the last edge of P_i for $i \in [k]$. Then precisely one of the following is true:*

(1) There is a set S of k edges of B such that $B - S$ contains no x–y path.

(2) There are $k+1$ edge-disjoint x–y paths $P'_1, \dots, P'_{k+1}$ such that $\vec{e}_i$ is the first edge of P'_i and $\vec{f}_i$ the last edge of some P'_j with $j \in [k+1]$.

The proof of Theorem 11.5.1 is inspired by the Böhme, Göring and Harant's proof [BGH01] of Menger's Theorem for directed graphs. Its technical finesse makes it possible to handle the complex structure of bidirected graphs.

Proof. The proof is by strong induction on the sum of the lengths of the paths P_i. We may assume without loss of generality that $\sigma(e,x) = -$ for any edge e incident with x, since changing these signs affects neither what counts as an x–y path nor the edge-cleanness of x. Similarly, we may assume without loss of generality that $\sigma(e,y) = -$ for any edge e incident with y. This implies that no trail and in particular no path in B contains x or y as an internal vertex; we use this fact repeatedly in the proof without mentioning it explicitly.

If the set $\{e_1, e_2, \dots, e_k\}$ is as in (1), then we are done; so suppose otherwise. Then there is an x–y path P_{k+1} containing no e_i. If P_{k+1} is edge-disjoint from all P_i with $i \in [k]$, then we achieved (2); so suppose not, and let $\vec{e}$ be the first edge of P_{k+1} which lies on some other P_i. Without loss of generality, we may assume that $e \in E(P_k)$. If $\overleftarrow{e}$ is in $\vec{E}(P_k)$, then $P_{k+1}\vec{e}P_k^-$ is an x–x trail, contradicting the edge-cleanness of x. So we must have $\vec{e} \in \vec{E}(P_k)$.

Let v be the startvertex of $\vec{e}$ and w its endvertex. We define R_1 as $P_k v$ and R_2 as $P_{k+1}v$. By definition, R_2 is edge-disjoint from all P_i with $i \leqslant k$, and it can be extended

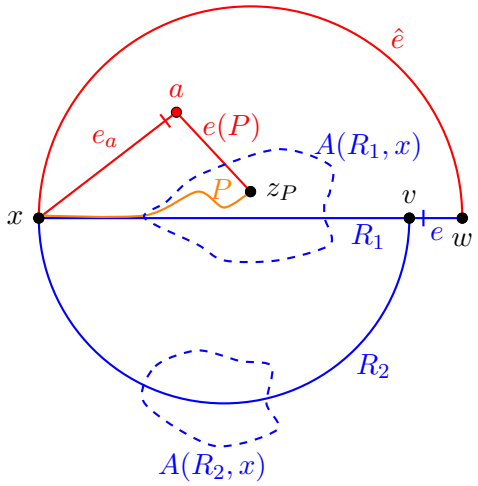

Figure 11.4: Construction of the auxiliary graph $\hat{B}$ in the proof of Theorem 11.5.1.

as a path by adding the edge e; these are the two properties of R_2 which we will use in the below argument. Note that R_1 has the same properties with respect to the paths $P_1, \dots, P_{k-1}$ and $R_2 v P_k$. We will exploit this symmetry repeatedly in what follows.

Let $A_i := A(R_i, x) \cup E(R_i)$. By Lemma 11.4.4, A_1 is disjoint from $E(R_2 v P_k)$ and A_2 is disjoint from $E(R_1 v P_k)$. Similarly, both A_1 and A_2 are disjoint from $E(P_i)$ for $i \in [k-1]$. To apply induction, we now construct a suitable auxiliary graph $\hat{B}$ using the sets A_1 and A_2.

We obtain $\hat{B}$ from B by modifying it in the following ways (see Figure 11.4):

- removing the edge e,
- adding a new edge $\hat{e}$ with endvertices x, w and $\sigma(\hat{e}, x) = -$ and $\sigma(\hat{e}, w) = \sigma(e, w)$,
- removing all edges in $A_1 \cup A_2$,
- adding a new vertex a,
- adding a new edge e_a from x to a with $\sigma(e_a, x) = -$ and $\sigma(e_a, a) = +$,
- adding a new edge g with endvertices a, z and $\sigma(a, g) = -$ and with $\sigma(z, g) = \alpha$ for each signed vertex (z, α) for which there is a nontrivial x–(z, α) path P, not just a trail, with $E(P) \subseteq A_i$ for some $i \in \{1, 2\}$. We will refer to g as $e(P)$ and to z as z_P.

We denote by $\vec{\hat{e}}$ and $\vec{e}_a$ the orientations of $\hat{e}$ and e_a, respectively, that point away from x. Furthermore, we write $\vec{e}(P)$ for the orientation of $e(P)$ that points away from a. For $i \in [k-1]$, let $\hat{P}_i := P_i$, and set $\hat{P}_k := \vec{\hat{e}} w P_k$. These k x–y paths $\hat{P}_i$ in $\hat{B}$ are edge-disjoint by construction. Note that $\hat{P}_k$ is shorter than P_k, since both R_i have length at least 1.

We now aim to apply the induction hypothesis to $\hat{B}$ and the paths $\hat{P}_i$ for $i \in [k]$. While the sum of the lengths of these paths is smaller than the one for $P_1, \dots, P_k$, we need to check that x is still edge-clean in $\hat{B}$.

So suppose for a contradiction that there is an x–x trail $\hat{Q}$ in $\hat{B}$. We aim to construct from $\hat{Q}$ an x–x trail Q in B, which then contradicts the edge-cleanness of x in B. For this, observe that the edge e_a is the only one with sign $+$ at a, and it is incident to x. So if $\hat{Q}$

contains any edge of the type $e(P)$, then $\hat{Q}$ also contains e_a directly before or after $e(P)$ and e_a must be the first or last edge of $\hat{Q}$. Similarly, the edge $\hat{e}$ can only occur as the first or last edge of $\hat{Q}$ since one of its endvertices is x.

So to transform $\hat{Q}$ into an x–x trail Q in B, we can focus on the first and last edge of $\hat{Q}$. Let us write $\vec{q}_1$ for the first edge of $\hat{Q}$ and $\overleftarrow{q}_2$ for its last edge. It then suffices to address the following cases; all remaining cases are symmetric by taking the inverse of $\hat{Q}$.

Case 1: q_1 and q_2 are both edges of B We can set $Q := \hat{Q}$.

Case 2: $q_1 = \hat{e}$ and q_2 is an edge of B We can set $Q := P_k w \hat{Q}$, since no edge of $P_k w$ is present in $\hat{B}$.

Case 3: $q_1 = e_a$ and q_2 is an edge of B The second edge of $\hat{Q}$ then has the form $\vec{e}(P)$ for some nontrivial path P with $E(P) \subseteq A_i$. So we can set $Q := P z_P \hat{Q}$, since no edge of A_i and hence P exists in $\hat{B}$.

Case 4: $q_1 = e_a$ and $q_2 = \hat{e}$ The second edge of $\hat{Q}$ then must be of the form $\vec{e}(P)$ for some nontrivial path P with $E(P) \subseteq A_i$. If $i = 1$, we can set $Q := P z_P \hat{Q} w \overleftarrow{e} v R_2^-$ for the same reasons as above; the case $i = 2$ then works analogously using R_1.

In all the above cases, we reach the desired contradiction, so we can conclude that x is still edge-clean in $\hat{B}$.

We can apply the induction hypothesis in $\hat{B}$ to the paths $\hat{P}_1, \dots \hat{P}_k$. This gives two possible outcomes, which we treat separately: either there is a set of edges satisfying (1) or there exists a set of $k+1$ paths satisfying (2).

Case (1) There is a set $\hat{S}$ of k edges of $\hat{B}$ such that $\hat{B} - \hat{S}$ contains no x–y path. Since the paths $\hat{P}_i$ are edge-disjoint, $\hat{S}$ consists of precisely one edge from each of these paths; in particular, $\hat{S}$ cannot contain e_a or any $e(P)$. We now set $S := \hat{S}$ if $\hat{e} \notin \hat{S}$ and let $S := (S \setminus \{\hat{e}\}) \cup \{e\}$ otherwise. Clearly, S is a set of k edges of B. We show that $B - S$ contains no x–y path, i.e. that S witnesses (1) to hold in B.

Suppose that there is an x–y path Q in $B - S$. To obtain the desired contradiction, we now construct from Q an x–y path $\hat{Q}$ in $\hat{B} - \hat{S}$. For this, we distinguish three cases, depending on whether Q meets edges of B that have been removed in $\hat{B}$.

Case 1.1: $E(Q)$ is disjoint from $A_1 \cup A_2 \cup \{e\}$ We can set $\hat{Q} := Q$.

Case 1.2: $\vec{e} \in \vec{E}(Q)$ and $E(wQ)$ is disjoint from $A_1 \cup A_2$ Since Q avoids S, we then have $e \notin S$, which implies $\hat{e} \notin \hat{S}$. Thus, we can set $\hat{Q} := x \vec{\hat{e}} w Q$.

Case 1.3: Otherwise Since $v \in V(A_1)$ is an endvertex of e and we are not in Case 1.1, Q meets $V(A_1 \cup A_2) \setminus \{x\}$. Let z be the last vertex of Q in this set. Since we are not in Case 1.2, we have $\overleftarrow{e} \in \vec{E}(Q)$ if e appears on Q; thus, v witnesses that some vertex in $V(A_1 \cup A_2) \setminus \{x\}$ then occurs on Q after e. Hence, the set $E(zQ)$ does not contain e.

We now assume that $z \in V(A_1) \setminus \{x\}$; the case of $z \in V(A_2) \setminus \{x\}$ is symmetric. Let α be the sign with which Q arrives at z. Applying Lemma 11.4.6 to the x–(z, α) path Qz in B, we find an x–(z, α) path P in A_1. Then we can set $\hat{Q} := x \vec{e}_a a \vec{e}(P) z Q$.

In each case, we find the desired contradiction, completing the proof of this case.

Case (2) There are $k+1$ edge-disjoint x–y paths $\hat{P}'_1, \hat{P}'_2, \dots, \hat{P}'_{k+1}$ in $\hat{B}$ such that the first edge of $\hat{P}'_i$ is $\vec{e}_i$ for $i \in [k-1]$, the first edge of $\hat{P}'_k$ is $\vec{\hat{e}}$, and for any $j \in [k]$ the last edge of $\hat{P}_j$ is the last edge of some $\hat{P}'_\ell$ with $\ell \in [k+1]$.
In this case, we will construct edge-disjoint x–y paths $P'_1, P'_2, \dots, P'_{k+1}$ in B such that the first edge of P'_i is $\vec{e}_i$ and $\vec{f}_i$ is the last edge of some P'_j with $j \in [k+1]$ for any $i \in [k]$.

For $i \in [k-1]$ the first edge of the path $\hat{P}'_i$ is $\vec{e}_i$, which is also its only edge incident to x. Thus, $\hat{P}'_i$ avoids the edge e_a, and therefore it also avoids any edge $e(P)$, as e_a is the only edge with sign $+$ at a and $e(P)$ has sign $-$ at a. Moreover, $\hat{P}'_i$ does not contain $\hat{e}$, the first edge of the path $\hat{P}_k$, from which $\hat{P}'_i$ is edge-disjoint. Thus, $\hat{P}'_i$ is a path in B, and we set $P'_i := \hat{P}'_i$ for $i \in [k-1]$. In particular, the first edge of any path P'_i with $i \in [k-1]$ is again $\vec{e}_i$.

It remains to define P'_k and P'_{k+1}. These two paths will only use edges of $\hat{P}'_k$ and $\hat{P}'_{k+1}$ as well as A_1 and A_2. Thus, they are edge-disjoint from the P'_i for $i \in [k-1]$; the edge-disjointness of $\hat{P}'_k$ and $\hat{P}'_{k+1}$ will follow directly from the construction.

We here again distinguish several cases, which depend on how $\hat{P}'_{k+1}$ meets those edges in $\hat{B}$ that do not exist in B.

Case 2.1: The first edge of $\hat{P}'_{k+1}$ is not $\vec{e}_a$ Then we can argue as for $\hat{P}'_i$ with $i \in [k-1]$ above, so $\hat{P}'_{k+1}$ is also a path in B. We thus set $P'_{k+1} := \hat{P}'_{k+1}$. Further, we let P'_k be the path obtained by applying Lemma 11.4.7 to $R_1 v \vec{e} w$ and $w\hat{P}'_k$ with respect to A_1 if $w \neq y$ and otherwise $P'_k := R_1 v \vec{e} w$.

Case 2.2 The first edge of $\hat{P}'_{k+1}$ is $\vec{e}_a$ and its second edge is $\vec{e}(P)$ for some nontrivial path P in A_2: We define P'_k as in Case 2.1. For P'_{k+1}, we take it as the path obtained by applying Lemma 11.4.7 to Pz_P and $z_P\hat{P}'_{k+1}$ with respect to A_2 if $z_P \neq y$, and we otherwise set $P'_{k+1} := Pz_P$.

Case 2.3 The first edge of $\hat{P}'_{k+1}$ is $\vec{e}_a$ and its second edge is $\vec{e}(P)$ for some nontrivial path P in A_1: We let P'_k be the path that we obtain by applying Lemma 11.4.7 to Pz_P and $z_P\hat{P}'_{k+1}$ with respect to A_1 if $z_P \neq y$ and otherwise $P'_k := Pz_P$. Moreover, we let P'_{k+1} be the path obtained by applying Lemma 11.4.7 to $R_2 v \vec{e} w$ and $w\hat{P}'_k$ with respect to A_2 if $w \neq y$ and otherwise $P'_{k+1} := R_2 v \vec{e} w$.

In all these cases, the resulting x–y paths P'_k and P'_{k+1} in B are edge-disjoint by their construction, and hence so are all $P'_1, \dots, P'_{k+1}$. As seen above, the first edge of any path P'_i with $i \in [k-1]$ is $\vec{e}_i$. We now verify that $\vec{e}_k$ is the first edge of P'_k. In Cases 2.1 and 2.2, the first edge of P'_k is the first edge of R_1, which is indeed $\vec{e}_k$. In Case 2.3, the first edge of P'_k is the first edge of P, which is a path in A_1 starting in x. By Corollary 11.4.5, this edge coincides indeed with the first edge of R_1, namely $\vec{e}_k$.

It remains to check that $\vec{f}_i$ is the last edge of some P'_j with $j \in [k+1]$ for any $i \in [k]$. Let $i \in [k]$ be arbitrary. If $\vec{f}_i = \vec{e}$, then $w = y$, and thus $\vec{f}_i$ is the last edge of either P'_k or P'_{k+1} in all the above cases. If $\vec{f}_i \neq \vec{e}$, then $\vec{f}_i$ is not only the last edge of P_i,

but also the last edge of $\hat{P}_i$ by the construction of $\hat{P}_i$. Thus, $\vec{f}_i$ is the last edge of some $\hat{P}'_\ell$ with $\ell \in [k+1]$. If $\hat{P}'_\ell$ is a path in B, then $\vec{f}_i$ is the last edge of $P'_\ell = \hat{P}'_\ell$. Otherwise, we have $\ell = k$ or $\ell = k+1$. Note that $\vec{f}_i \neq \vec{e}$ implies $w \neq y$. Further, since f_i is an edge of B, the edge f_i cannot occur as any $e(P)$. Hence, $\hat{P}'_\ell$ is not of the form $x\vec{e}_a a\vec{e}(P)z_P$ for some nontrivial path P in A_1 or A_2, and so $z_P \neq y$. So in the construction of $\hat{P}'_\ell$ in the above Cases 2.1 to 2.3, we applied Lemma 11.4.7, and the second path in this application was a nontrivial path of the form $w\hat{P}'_\ell$ or $z_P\hat{P}'_\ell$. Thus, $\vec{f}_i$ is the last edge of either P'_k or P'_{k+1}.

This completes the proof. □

Theorem 11.5.1 implies Menger's Theorem for edge-disjoint paths in both undirected and directed graphs. To show this, we may regard any graph as a directed graph by replacing each edge e with the two directed edges $\vec{e}$ and $\overleftarrow{e}$ (see [BG09] for an in-depth explanation). Similarly, we may regard any directed graph as a bidirected graph by viewing a directed edge $\vec{e}$ from x to y as an edge e with endvertices x and y and signs $-$ at x and $+$ at y.

We then obtain the desired version of Menger's Theorem for vertices x and y in a directed graph D by considering the directed graph D' obtained from D by deleting all edges that point towards x. Note that x is edge-clean in D' since no trail ends in x. If there is no x–y path in $D' - S$ for some $S \subseteq V(D')$, then there is no x–y path in $D - S$. Furthermore, any x–y path in D' is also an x–y path in D.

11.6 Menger's theorem for vertex-disjoint paths

In this section we deduce the vertex-disjoint version of Menger's Theorem for bidirected graphs, Theorem 31, from the above shown edge-version, Theorem 11.5.1. Theorem 31 will be formulated in terms of vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths between two sets $\mathcal{X}$ and $\mathcal{Y}$ of signed vertices of a bidirected graph B. Analogously to the edge-version where the startvertex x of the paths had to be edge-clean, we have to require the set $\mathcal{X}$ of signed vertices to be *clean*, which is defined as follows:

Definition 11.6.1. A set $\mathcal{X}$ of signed vertices of a bidirected graph B is *clean* if B contains no nontrivial path starting and ending in $\mathcal{X}$.

We remark that a clean set $\mathcal{X}$ of signed vertices may contain both $(x,+)$ and $(x,-)$ for a vertex x.

Let us now recall the vertex-disjoint version of Menger's Theorem for bidirected graphs from the introduction:

Theorem 31. *Let $\mathcal{X}$ and $\mathcal{Y}$ be sets of signed vertices of a bidirected graph B, and suppose that $\mathcal{X}$ is clean. Let $P_1, \dots, P_k$ be vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths in B where P_i starts in $v_i \in V(\mathcal{X})$ and ends in $w_i \in V(\mathcal{Y})$ for $i \in [k]$. Then precisely one of the following is true:*

(1) There is a set S of k vertices of B such that $B - S$ contains no $\mathcal{X}$–$\mathcal{Y}$ path.

(2) There are $k+1$ vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths $P'_1, \dots, P'_{k+1}$ in B such that P'_i starts in v_i and some P'_j with $j \in [k+1]$ ends in w_i.

Note that in Theorem 31 we only require $\mathcal{X}$ to be clean, but not $\mathcal{Y}$ – just as we only require x (and not y) to be edge-clean in Theorem 11.5.1; indeed, we do not need any assumptions on $\mathcal{Y}$. We further note that Theorem 31 cannot be strengthened by fixing the signed startvertices of the paths P_i rather than just their startvertices, see Figure 11.5 for a counterexample.

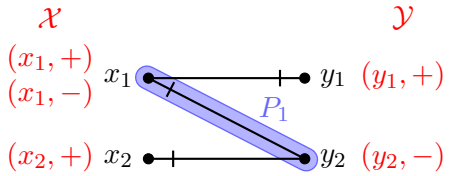

Figure 11.5: This bidirected graph contains two disjoint $\mathcal{X}$–$\mathcal{Y}$ paths, namely the paths $x_2(x_2y_2, x_2, y_2)y_2$ and $x_1(x_1y_1, x_1, y_1)y_1$. However, while the $\mathcal{X}$–$\mathcal{Y}$ path $P_1 := x_1(x_1y_2, x_1, y_2)y_2$ starts in $(x_1,+) \in \mathcal{X}$, the graph contains no two disjoint $\mathcal{X}$–$\mathcal{Y}$ paths such that one of them starts in $(x_1,+)$.

Towards a proof of Theorem 31, let us first show that we may assume the $\mathcal{X}$–$\mathcal{Y}$ paths in B to be the only paths starting in $\mathcal{X}$ and ending in $\mathcal{Y}$, respectively. More precisely, by removing some edges from B and changing some signs at vertices of $V(\mathcal{X}) \cup V(\mathcal{Y})$, all nontrivial paths starting in $\mathcal{X}$ and ending in $\mathcal{Y}$ are indeed $\mathcal{X}$–$\mathcal{Y}$ paths, i.e. they are internally disjoint from $V(\mathcal{X}) \cup V(\mathcal{Y})$ and contain no trivial $\mathcal{X}$–$\mathcal{Y}$ paths:

Proposition 11.6.2. *Let B be a bidirected graph and $\mathcal{X}, \mathcal{Y}$ sets of signed vertices in B. Then there exists a bidirected graph B' with $V(B') = V(B)$, $E(B') \subseteq E(B)$ and a set of signed vertices $\mathcal{X}' \subseteq \mathcal{X}$ in B' such that*

(1) any $\mathcal{X}'$–$\mathcal{Y}$ path in B' is an $\mathcal{X}$–$\mathcal{Y}$ path in B and vice versa,

(2) any path in B' starting in $\mathcal{X}'$ and ending in $\mathcal{Y}$ is an $\mathcal{X}'$–$\mathcal{Y}$ path, and

(3) there is no trivial $\mathcal{X}'$–$\mathcal{X}'$ path in B'.

Moreover, if $\mathcal{X}$ is clean in B, then $\mathcal{X}'$ is clean in B'.

Note that by (1), Theorem 31, or more generally any version of Menger's theorem concerning vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths, holds for B, $\mathcal{X}$ and $\mathcal{Y}$ if and only if it holds for B', $\mathcal{X}'$ and $\mathcal{Y}$.

Proof. We construct B' and $\mathcal{X}'$ from B and $\mathcal{X}$ by performing the following modifications for every $v \in V(\mathcal{X}) \cup V(\mathcal{Y})$ simultaneously:

(a) If there is $\alpha \in \{+,-\}$ with $(v,\alpha) \notin \mathcal{X} \cup \mathcal{Y}$, we delete all edges incident to v that have sign α at v.

(b) If $(v,+)$ and $(v,-)$ are either both contained in $\mathcal{X} \setminus \mathcal{Y}$ or both contained in $\mathcal{Y} \setminus \mathcal{X}$, we change the signs of all edges incident to v at v to $+$. Furthermore, if $(v,-) \in \mathcal{X}$, then we remove $(v,-)$ from $\mathcal{X}$.

(c) If there is $\alpha \in \{+,-\}$ such that $(v,\alpha) \in \mathcal{X}$ and $(v,-\alpha) \in \mathcal{Y}$, we isolate v, i.e. we delete all edges incident with v. If $(v,-\alpha) \in \mathcal{X}$, then we also delete $(v,-\alpha)$ from $\mathcal{X}$.

With this construction, we in particular obtain the following properties of B':

(i) $V(\mathcal{X}) \cup V(\mathcal{Y}) = V(\mathcal{X}') \cup V(\mathcal{Y})$, and all edges incident to $v \in V(\mathcal{X}') \cup V(\mathcal{Y})$ have the same sign at v in B'.

(ii) If an edge of B' is incident to a vertex $w \notin V(\mathcal{X}') \cup V(\mathcal{Y})$, then it has the same sign at w in B' and in B.

(iii) If a vertex v forms a trivial $\mathcal{X}$–$\mathcal{Y}$ path in B, then it is an isolated vertex in B' and forms a trivial $\mathcal{X}'$–$\mathcal{Y}$ path in B'.

We begin by proving (1). For trivial paths, the statement is true by (iii) and since $\mathcal{X}' \subseteq \mathcal{X}$. So let P be a nontrivial $\mathcal{X}$–$\mathcal{Y}$ path in B. Then neither its start- nor its endvertex forms a trivial $\mathcal{X}'$–$\mathcal{Y}$ path in B' since they each do not form a trivial $\mathcal{X}$–$\mathcal{Y}$ path in B. In particular, we did not apply modification (c) to the start- or endvertex of P. As all internal vertices of P are disjoint to $V(\mathcal{X}) \cup V(\mathcal{Y})$, we did not apply any modification to these vertices. This implies that all edges of P are edges in B'. By (ii) P is thus also a path in B'. Clearly, it is internally disjoint to $V(\mathcal{X}') \cup V(\mathcal{Y})$. Finally, since we applied either modification (a) or (b) to its start- and endvertex, the path P in B' starts in $\mathcal{X}'$ and ends in $\mathcal{Y}$.

Let P' be a nontrivial $\mathcal{X}'$–$\mathcal{Y}$ path in B'. Then the path P' is, by (i) internally disjoint to $V(\mathcal{X}) \cup V(\mathcal{Y}) = V(\mathcal{X}') \cup V(\mathcal{Y})$ and thus also a path in B by (ii). As (1) holds for trivial paths, neither its start- nor its endvertex forms a trivial $\mathcal{X}$–$\mathcal{Y}$ path in B. In particular, we did not apply modification (c) to its start- or endvertex. Thus, we applied either modification (a) or (b) to its start- and endvertex. This implies that P' in B starts in $\mathcal{X}$ and ends in $\mathcal{Y}$, which completes the proof of (1).

For (2), consider a nontrivial path P' in B' that starts in $\mathcal{X}'$ and ends in $\mathcal{Y}$. To show that P' is an $\mathcal{X}'$–$\mathcal{Y}$ path, we need to verify two properties: First, no internal vertex of P' is contained in $V(\mathcal{X}') \cup V(\mathcal{Y})$, which is true by (i). Second, P' contains no trivial $\mathcal{X}'$–$\mathcal{Y}$ path. Indeed, if a vertex $v \in P'$ is a trivial $\mathcal{X}'$–$\mathcal{Y}$ path in B', then it is also a trivial $\mathcal{X}$–$\mathcal{Y}$ path in B by (1) and hence isolated by (iii), which contradicts that P' is nontrivial.

For (3), note that the modifications (a) to (c) ensure that $(v,\alpha) \in \mathcal{X}'$ implies $(v,-\alpha) \notin \mathcal{X}'$.

For the 'moreover'-part, assume that $\mathcal{X}$ is clean in B and suppose for a contradiction that x is not edge-clean in B'. Then there exists a nontrivial path P in B' that starts and ends in $\mathcal{X}'$. By (i), no internal vertex of P is contained in $V(\mathcal{X}) \cup V(\mathcal{Y})$. Furthermore, by (ii), P still forms a path in B. Thus, P is a nontrivial path in B starting and ending in $\mathcal{X}$, which contradicts the cleanness of $\mathcal{X}$ in B. □

From now on, let B be a bidirected graph and let $\mathcal{X}$ and $\mathcal{Y}$ be sets of signed vertices of B. By Proposition 11.6.2 we may assume that B satisfies (2) and (3). To prove Theorem 31 for B, $\mathcal{X}$ and $\mathcal{Y}$, we now carefully construct an auxiliary graph $\hat{B}$ with two specified vertices x and y, such that Theorem 11.5.1, our edge-disjoint version of Menger's theorem

in bidirected graphs, applied to $\hat{B}$, x and y allows us to deduce Theorem 31 for B, $\mathcal{X}$ and $\mathcal{Y}$.

The construction of $\hat{B}$, however, turns out to be even more general: It provides a framework to deduce a vertex-version of Menger's theorem in bidirected graphs from an edge-version. In particular, our construction of $\hat{B}$ is independent of the cleanness assumption on $\mathcal{X}$ as in Theorem 31, and the cleanness of $\mathcal{X}$ in B will only be used to make Theorem 11.5.1 applicable to x in $\hat{B}$ in that it implies x to be edge-clean.

In order to make the above approach work, we construct $\hat{B}$, x and y such that the vertices x and y of $\hat{B}$ represent the sets $\mathcal{X}$ and $\mathcal{Y}$ of signed vertices of B, respectively. They do so in such a way that any set of vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths in B correspond to edge-disjoint x–y paths in $\hat{B}$ and vice versa.

To show that a vertex-version of Menger's theorem for B, $\mathcal{X}$ and $\mathcal{Y}$ can be deduced from an edge-version of Menger's theorem for $\hat{B}$, x and y, we then verify two properties of our construction. We first show that an edge-separator for x and y in $\hat{B}$ corresponds to a vertex-separator for $\mathcal{X}$ and $\mathcal{Y}$ in B of at most the same size. Secondly, we consider vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths $P_1, \dots, P_k$ in B. By our construction, they correspond to edge-disjoint x–y paths $\hat{P}_1, \dots, \hat{P}_k$ in $\hat{B}$. We then show that $k+1$ edge-disjoint x–y paths in $\hat{B}$ such that all but one start in the same edges as $\hat{P}_1, \dots, \hat{P}_k$ and all but one end in the same edges as $\hat{P}_1, \dots, \hat{P}_k$ correspond to $k+1$ vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths such that all but one start in the same vertices as $P_1, \dots, P_k$ and all but one end in the same vertices as $P_1, \dots, P_k$.

Finally, we show that our construction indeed allows us to apply Theorem 11.5.1 to $\hat{B}$, x and y under the assumptions of Theorem 31 on B, $\mathcal{X}$ and $\mathcal{Y}$. More precisely, we prove that the cleanness of $\mathcal{X}$ in B implies that x is edge-clean in $\hat{B}$.

Let us now construct $\hat{B} = (\hat{G}, \hat{\sigma})$ from $B = (G, \sigma)$ (see Figure 11.6 for an illustration of this construction). For every vertex $v \in B$, the vertex set of $\hat{B}$ contains two vertices v^+ and v^- which are connected in $\hat{B}$ by an edge with sign $+$ at v^- and sign $-$ at v^+. An edge $e \in E(B)$ which is incident to u and v in B then transfers to an edge $\hat{e} \in E(\hat{B})$ that is incident to $u^{\sigma(u,e)}$ and $v^{\sigma(v,e)}$. The half-edges of $\hat{e}$ have the same signs as the respective half-edges of e, that is $\hat{\sigma}(u^{\sigma(u,e)}, \hat{e}) := \sigma(u, e)$ and $\hat{\sigma}(v^{\sigma(v,e)}, \hat{e}) := \sigma(v, e)$.

We finally add two new vertices x and y to $V(\hat{B})$, which shall represent the sets $\mathcal{X}$ and $\mathcal{Y}$, respectively, as follows. The vertex x is adjacent to those $v^\alpha \in V(\hat{B})$ with $(v, -\alpha) \in \mathcal{X}$, and the respective edge in $\hat{B}$ has sign α at v^α and sign $-$ at x. Analogously, the vertex y is adjacent to every $v^\beta \in V(\hat{B})$ with $(v, -\beta) \in \mathcal{Y}$, and the respective edge has sign β at v^β and sign $-$ at y.

Observe that with this construction of $\hat{B} = (\hat{G}, \hat{\sigma})$, the graph $\hat{G}$ has no parallel edges. For notational simplicity, we will thus identify an (oriented) edge in $\hat{B}$ with the (ordered) pair of its endvertices in what follows.

Now we turn our attention to the relation between paths in B and paths in $\hat{B}$. Let $P = v_1 \vec{e}_1 v_2 \dots v_{n-1} \vec{e}_{n-1} v_n$ be any $\mathcal{X}$–$\mathcal{Y}$ path in B, and write $\alpha_\ell := \sigma(v_\ell, e_\ell)$ and $\beta_\ell := \sigma(v_{\ell+1}, e_\ell)$ for all $\ell \in [n-1]$. Since P is an $\mathcal{X}$–$\mathcal{Y}$ path in B, we have $\beta_\ell = -\alpha_{\ell+1}$ for all $\ell \in [n-2]$ as well as $(v_1, \alpha_1) \in \mathcal{X}$ and $(v_n, \beta_{n-1}) \in \mathcal{Y}$. Thus, the construction of $\hat{B}$ guarantees that

$$\begin{aligned}\hat{P} :=& x(x, v_1^{-\alpha_1}) v_1^{-\alpha_1} (v_1^{-\alpha_1}, v_1^{\alpha_1}) v_1^{\alpha_1} (v_1^{\alpha_1}, v_2^{\beta_1}) v_2^{\beta_1} \dots \\ & \dots v_{n-1}^{\alpha_{n-1}} (v_{n-1}^{\alpha_{n-1}}, v_n^{\beta_{n-1}}) v_n^{\beta_{n-1}} (v_n^{\beta_{n-1}}, v_n^{-\beta_{n-1}}) v_n^{-\beta_{n-1}} (v_n^{-\beta_{n-1}}, y) y\end{aligned}$$

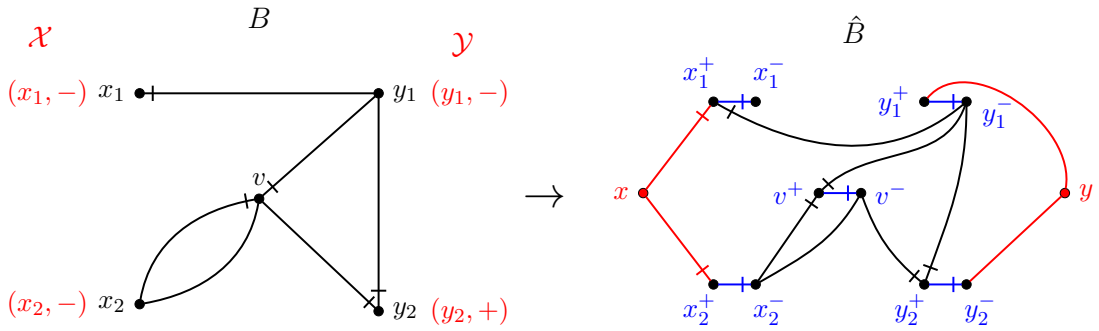

Figure 11.6: The construction of $\hat{B}$ from the bidirected graph B with specified sets $\mathcal{X}$ and $\mathcal{Y}$ of signed vertices as in Section 11.6.

is an x–y path in $\hat{B}$. Moreover, if $\mathcal{X}$–$\mathcal{Y}$ paths $P_1, \dots, P_k$ are vertex-disjoint, then the corresponding x–y paths $\hat{P}_1, \dots, \hat{P}_k$ in $\hat{B}$ are internally vertex-disjoint and in particular edge-disjoint.

With this transfer of $\mathcal{X}$–$\mathcal{Y}$ paths in B to x–y paths in $\hat{B}$ at hand, it is easy to show that an edge-separator of x and y in $\hat{B}$ yields a vertex-separator of $\mathcal{X}$ and $\mathcal{Y}$ in B of at most the same size:

Lemma 11.6.3. *If there is a set $\hat{S} \subseteq E(\hat{B})$ of size at most k such that there is no x–y path in $\hat{B} - \hat{S}$, then there is a set $S \subseteq V(B)$ of size at most k such that there is no $\mathcal{X}$–$\mathcal{Y}$ path in $B - S$.*

Proof. Let $\hat{S}$ be a set of at most k edges in $\hat{B}$ such that there is no x–y path in $\hat{B} - \hat{S}$. We define a set S of k vertices of B as follows:

- ◇ if $\hat{e} \in \hat{S}$ has the form $v^\alpha v^{-\alpha}$, then we add v to S,
- ◇ if $\hat{e} \in \hat{S}$ has the form $v^\alpha w^\beta$, then we arbitrarily add one of v and w to S, and
- ◇ if $\hat{e} \in \hat{S}$ has the form $v^\alpha x$ or $v^\alpha y$, then we add v to S.

Then there is no $\mathcal{X}$–$\mathcal{Y}$ path P in $B - S$, as the corresponding path $\hat{P}$ in $\hat{B}$ would clearly avoid $\hat{S}$. □

We now turn to the transfer of x–y trails in $\hat{B}$ to $\mathcal{X}$–$\mathcal{Y}$ paths in B. For this, let us first observe that such trails admit a very particular structure due to the construction of $\hat{B}$:

Proposition 11.6.4. *Let T be either an x–y trail or a nontrivial x–x trail in $\hat{B}$. Then*

(1) every internal vertex of T has the form v^α for some $v \in V(B)$ and $\alpha \in \{+,-\}$,

(2) every internal vertex of T is met by T exactly once, and

(3) precisely every second edge of T is of the form $(v^\alpha, v^{-\alpha})$.

Proof. All edges incident to x or y have sign $-$ at x or y, so every internal vertex of a trail in $\hat{B}$ has the form v^α for some $v \in V(B)$ and $\alpha \in \{+,-\}$, which shows (1). For every vertex of this form v^α (i.e. every vertex of $\hat{B}$ other than x and y), there exists precisely one edge in $\hat{B}$ which has sign $-\alpha$ at v^α, namely the edge $v^\alpha v^{-\alpha}$. This implies that T uses

the edge $(v^\alpha, v^{-\alpha})$ either immediately before or after meeting v^α. In particular, (3) holds. Since v^α is neither the start- nor the endvertex of T, it is met by T exactly once, which proves (2). □

Now let $\hat{T} := v_1\vec{\hat{e}}_1 v_2 \vec{\hat{e}}_2 \dots \vec{\hat{e}}_{n-2} v_{n-1} \vec{\hat{e}}_{n-1} v_n$ be either an x–y trail or a nontrivial x–x trail in $\hat{B}$. By Proposition 11.6.4 (1) and (2), the trail $\hat{T}$ meets all vertices of $\hat{B}$ but x at most once. Thus, its subtrail $\hat{T}' := \vec{\hat{e}}_2 \hat{T} \vec{\hat{e}}_{n-2}$ is a path.

Now replace every subtrail of $\hat{T}'$ that has the form $v^\alpha(v^\alpha, v^{-\alpha})v^{-\alpha}$ by the vertex $v \in B$; note in particular that both $v_2\vec{\hat{e}}_2 v_3$ and $v_{n-2}\vec{\hat{e}}_{n-2}v_{n-1}$ have this form. Moreover, replace any edge of $\hat{T}'$ that has the form (v^α, w^β) for distinct $v, w \in V(B)$ by the oriented edge in B that has sign α at v and sign β at w (such an edge exists by the construction of $\hat{B}$). Proposition 11.6.4 (3) then guarantees that the constructed P is indeed a path in B.

We show that P starts in $\mathcal{X}$; the case of P ending in $\mathcal{X}$ or in $\mathcal{Y}$ is symmetrical. The only neighbours of x in $\hat{B}$ are the vertices v^α with $(v, -\alpha) \in \mathcal{X}$. Proposition 11.6.4 (3) then implies that if $v_2 = v^\alpha$, then $\vec{\hat{e}}_2 = (v^\alpha, v^{-\alpha})$ and so $\hat{e}_3$ has sign $-\alpha = -\hat{\sigma}(v_3, \hat{e}_2)$ at $v_3 = v^{-\alpha}$. Thus, our construction of P guarantees that P starts in $(v, -\alpha) \in \mathcal{X}$, as desired.

So if $\hat{T}$ is an x–y trail, then P starts in $\mathcal{X}$ and ends in $\mathcal{Y}$ and it is indeed an $\mathcal{X}$–$\mathcal{Y}$ path in B, since we assumed B to satisfy (2) in Proposition 11.6.2. And if $\hat{T}$ is a nontrivial x–x trail, then P starts and ends in $\mathcal{X}$ and it is nontrivial, since B also satisfies (3) by assumption.

This transfer of x–y trails from $\hat{B}$ to $\mathcal{X}$–$\mathcal{Y}$ paths in B now interacts with the transfer from B to $\hat{B}$ precisely in the desired way:

Lemma 11.6.5. *Let $P_1, \dots, P_k$ be $\mathcal{X}$–$\mathcal{Y}$ paths in B, and let $\hat{P}_1, \dots, \hat{P}_k$ be the corresponding x–y paths in $\hat{B}$. Further, let $\hat{P}'_1, \dots, \hat{P}'_{k+1}$ be $k+1$ edge-disjoint x–y paths in $\hat{B} = (\hat{G}, \hat{\sigma})$ where $\hat{P}'_i$ starts in the same edge as $\hat{P}_i$ and some $\hat{P}'_j$ with $j \in [k+1]$ ends in the same edge as $\hat{P}_i$ for any $i \in [k]$. Then there are $k+1$ vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths $P'_1, \dots, P'_{k+1}$ in B where P'_i starts in the same vertex as P_i and some P'_j with $j \in [k+1]$ ends in the same vertex as P_i for any $i \in [k]$.*

Proof. Let $P'_1, \dots, P'_{k+1}$ be the $\mathcal{X}$–$\mathcal{Y}$ paths in B obtained from $\hat{P}'_1, \dots, \hat{P}'_{k+1}$. Let $i \in [k]$ be arbitrary, and let v^α be the endvertex of the first edge of $\hat{P}'_i$. By construction, the path P'_i starts in the vertex v. Now our assumption on $\hat{P}'_i$ implies that the path $\hat{P}'_i$ starts in the same edge as $\hat{P}_i$. By the construction of $\hat{P}_i$ from P_i, the endvertex v^α of the first edge of $\hat{P}_i$ has the property that v is the startvertex of P_i. Thus, P_i and P'_i start in the same vertex, as desired.

Let $j \in [k+1]$ be such that $\hat{P}'_j$ ends in the same edge as $\hat{P}_i$. By a similar argument as for the startvertex, the path P'_j ends in the same vertex as P_i.

It remains to prove that the paths $P'_1, \dots, P'_{k+1}$ are vertex-disjoint: Proposition 11.6.4 (3) together with our construction yields that if two distinct P'_i and P'_j would contain the same vertex v, then both $\hat{P}'_i$ and $\hat{P}'_j$ would have to contain the edge $v^\alpha v^{-\alpha}$, which contradicts that they are edge-disjoint. □

With Lemma 11.6.3 and Lemma 11.6.5, we have now shown that the construction of $\hat{B}$ indeed provides a general framework to deduce a vertex-version of Menger's theorem

in bidirected graphs form an edge-version. We can thus finally deduce Theorem 31 from Theorem 11.5.1 by checking that x is edge-clean in $\hat{B}$ if $\mathcal{X}$ is clean in B:

Proof of Theorem 31. By Proposition 11.6.2, we may assume that B, $\mathcal{X}$ and $\mathcal{Y}$ satisfy (2) and (3), and this does not affect the cleanness of $\mathcal{X}$ in B. Let $\hat{B}$, x and y be constructed from B, $\mathcal{X}$ and $\mathcal{Y}$ as above, and suppose for a contradiction that x is not edge-clean, i.e. that there exists a nontrivial x–x trail $\hat{T}$ in $\hat{B}$. As shown above, the trail $\hat{T}$ then gives rise to a nontrivial path in B that starts and ends in $\mathcal{X}$, which contradicts the cleanness of $\mathcal{X}$ in B. Thus, we can apply Theorem 11.5.1 to $\hat{B}$, x, y and the k edge-disjoint x–y paths $\hat{P}_1, \dots, \hat{P}_k$ in $\hat{B}$ corresponding to the k vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths $P_1, \dots, P_k$ in B. Depending on the outcome of Theorem 11.5.1, Lemma 11.6.3 or Lemma 11.6.5 then complete the proof. □

We remark that it is easy to deduce a version of Theorem 31 in which an $\mathcal{X}$–$\mathcal{Y}$ path does not have to be internally disjoint from $V(\mathcal{X}) \cup V(\mathcal{Y})$ and may contain trivial $\mathcal{X}$–$\mathcal{Y}$ paths: Apply Theorem 31 to the bidirected graph that is obtained by adding a new vertex for every $v \in V(\mathcal{X})$ which is joined to v by an edge with sign $-\alpha$ at v and sign α at the newly added vertex for any $(v, \alpha) \in \mathcal{X}$, and by adding vertices and edges for $v \in V(\mathcal{Y})$ in the same way. Then take $\mathcal{X}'$ and $\mathcal{Y}'$ as the respective sets of newly added vertices.

Theorem 31 implies Menger's Theorem for vertex-disjoint paths in both graphs and directed graphs. As explained at the end of Section 11.5, we may regard any undirected or directed graph as a bidirected graph. We then obtain the desired version of Menger's Theorem for vertex sets X and Y in a directed graph D by considering the sets $\mathcal{X} := \{(x, -) \mid x \in X\}$ and $\mathcal{Y} = \{(y, +) \mid y \in Y\}$ of signed vertices. Note that $\mathcal{X}$ is indeed clean since every path in the directed graph D which ends in X does so in $(x, +)$ for some $x \in X$ and hence not in $\mathcal{X}$.

11.7 Polynomial time algorithm

In this section we show that the proof of the edge-version of Menger's Theorem for bidirected graphs, Theorem 11.5.1, can be turned into a polynomial time algorithm. More precisely, we show the following statement:

Theorem 11.7.1. *There exists a polynomial time algorithm that, given distinct vertices x and y of a bidirected graph B and k edge-disjoint x–y paths $P_1, \dots, P_k$ in B where P_i starts in $\vec{e}_i$ and ends in $\vec{f}_i$ for $i \in [k]$, finds either a set S of k edges of B such that $B - S$ does not contain any x–y path or $k+1$ edge-disjoint x–y paths $P'_1, \dots, P'_{k+1}$ in B such that P'_i starts in $\vec{e}_i$ and some P'_j with $j \in [k+1]$ ends in $\vec{f}_i$ for any $i \in [k]$.*

As a consequence, there exists a polynomial time algorithm for Theorem 31 since the transfer process introduced in Section 11.6 can be computed in polynomial time.

For the proof of Theorem 11.7.1 note that the proof of Theorem 11.5.1 uses a recursive procedure by induction on the sum of the lengths of the paths P_i; it is thus enough to show that each recursion step runs in polynomial time as the number of recursion steps is bounded by $|E(B)|$.

To show that each recursion step runs in polynomial time, we will in particular need polynomial time algorithms to find paths and trails between two fixed signed vertices as

well as the appendage of a given path. One key ingredient to this is Edmonds' celebrated Blossom Algorithm:

Theorem 11.7.2 ([E65])**.** *There exists a polynomial time algorithm which, given a graph G, either finds a perfect matching in G or determines that G has no perfect matching.*

We first show that a path between two signed vertices can be found in polynomial time:

Lemma 11.7.3. *There exists a polynomial time algorithm that, given a bidirected graph B and signed vertices (x,α) and (y,β) of B with $x \neq y$, either finds an (x,α)–(y,β) path in B or determines that no such path exists.*

Proof. We first describe an algorithm which runs in polynomial time and then prove its correctness.

The algorithm starts by checking if there exists an edge in B with sign α at x and sign β at y, which then forms our desired path. If no such edge exists, then the algorithm considers the graph $B' := B - \{x,y\}$ and computes the set $\mathcal{X}'$ of all signed vertices (v,γ) of B' for which there exists an (x,α)–(v,γ) edge and the set $\mathcal{Y}'$ which is defined analogously with respect to y. As in Section 11.6, the algorithm then constructs the auxiliary graph $\hat{B}' = (\hat{G}', \hat{\sigma}')$ from B', $\mathcal{X}'$ and $\mathcal{Y}'$; let x' and y' be the two vertices of $\hat{B}'$ representing $\mathcal{X}'$ and $\mathcal{Y}'$, respectively. Note that the construction of $\hat{B}'$ can be done in polynomial time.

The algorithm then uses the polynomial time algorithm of Theorem 11.7.2 to construct a perfect matching of $\hat{G}'$ if possible. If there is a perfect matching, then the algorithm computes an (x,α)–(y,β) path in B from it in polynomial time (see below). Otherwise, the algorithm returns that there is no (x,α)–(y,β) path in B. This concludes our description of the polynomial time algorithm.

For the proof of its correctness, let $\hat{e}_v$ denote the unique edge of $\hat{G}'$ between v^+ and v^- for every vertex $v \in V(B')$. Then $M := \{\hat{e}_v \mid v \in V(B')\}$ is a matching in $\hat{G}'$, and every vertex of $\hat{G}'$ except from x' and y' is incident to an edge in M.

If there is a perfect matching M' in $\hat{G}'$, then let H' be the subgraph of $\hat{G}'$ with edge set $M \cup M'$. By the definition of M, the component of H' containing x' is an x'–y' path $\hat{P}$. The construction of $\hat{B}'$ together with the definition of M implies that such a path $\hat{P}$ in $\hat{G}'$ induces an x'–y' path $\hat{P}'$ in $\hat{B}'$. As in the proof of Theorem 31, this path $\hat{P}'$ then induces an $\mathcal{X}'$–$\mathcal{Y}'$ path P' in B'. The definition of $\mathcal{X}'$ and $\mathcal{Y}'$ then allows us to extend the path P' to an (x,α)–(y,β) path in B.

Conversely, we show that if there exists an (x,α)–(y,β) path in B, then there is a perfect matching in $\hat{G}$. Let $P = x\vec{e}_1 v_1 \dots v_{\ell-1}\vec{e}_\ell y$ be an (x,α)–(y,β) path in B. Then $P' = v_1 \dots v_{\ell-1}$ is an $\mathcal{X}'$–$\mathcal{Y}'$ path in B' by the definition of $\mathcal{X}'$ and $\mathcal{Y}'$. As in the proof of Theorem 31, the path P' induces an x'–y' path $\hat{P}'$ in $\hat{B}'$. The respective path $\hat{Q}'$ in $\hat{G}'$ alternates between edges in M and not in M. Note that all the vertices of $\hat{G}'$ but the startvertex and endvertex of $\hat{Q}'$ are incident to edges in M. This means that no vertex of $\hat{Q}'$ is incident to an edge in $M \setminus E(\hat{Q}')$. It follows that the symmetric difference $(M \setminus E(\hat{Q}')) \cup (E(\hat{Q}') \setminus M)$ is a perfect matching in $\hat{G}'$, as desired. □

Next, we extend Lemma 11.7.3 from paths to trails. Our main tool for this is the line graph of a bidirected graph which we introduce below. For this we need one more definition:

An *orientation of a bidirected graph* B is a map $\nu : E(B) \to \overleftrightarrow{E(B)}$ such that $\nu(e) \in \{\vec{e}, \overleftarrow{e}\}$ for any $e \in E(B)$.

Definition 11.7.4. Given a bidirected graph $B = (G, \sigma)$ and an orientation ν of B, we define the *line graph of B with respect to ν*, denoted as $L := L_\nu = (H, \tau)$, as follows. The graph H has vertex set $E(G)$ and contains an edge $\{e_1, e_2\} \in E(H)$ with label v for any two distinct edges $e_1, e_2 \in E(G)$ and a vertex $v \in e_1 \cap e_2$ satisfying $\sigma(v, e_1) = -\sigma(v, e_2)$. The signing τ of the half-edges $\mathbf{E}(H)$ assigns to an edge $\{e_1, e_2\} \in E(H)$ with label v the sign

$$\tau(e_i, \{e_1, e_2\}) := \begin{cases} + & \text{if } v \text{ is endvertex of } \nu(e_i), \\ - & \text{if } v \text{ is startvertex of } \nu(e_i), \end{cases}$$

for $i = 1, 2$.

Fix an arbitrary orientation of a bidirected graph B, and let L be the corresponding line graph of B. Just as for line graphs of (un)directed graphs, the above definition of L implies that walks in B induce walks in L and vice versa. Indeed, a walk $v_0\vec{e}_1 v_1 \dots v_{n-1}\vec{e}_n v_n$ in B induces the walk $e_1\vec{a}_1 e_2 \dots e_{n-1}\vec{a}_{n-1} e_n$ in L where $\vec{a}_i$ is the directed edge from e_i to e_{i+1} with label v_i for any $i \in [n-1]$. Vice versa, consider a walk $W = e_1\vec{a}_1 e_2 \dots e_{n-1}\vec{a}_{n-1} e_n$ in L which has label v_i at the edge a_i for any $i \in [n-1]$, and let $v_0 \in e_1 \setminus \{v_1\}$ and $v_n \in e_n \setminus \{v_{n-1}\}$. Writing $\vec{e}_i = (e_i, v_{i-1}, v_i)$ for $1 \leqslant i \leqslant n$, the walk W in L then induces the walk $v_0\vec{e}_1 v_1 \dots \vec{e}_n v_n$ in B.

The above correspondences between walks in B and walks in L imply in particular that trails in B induce paths in L and paths in L induce trails in B. We can thus compute trails in B in polynomial time by applying Lemma 11.7.3 to a suitable line graph of B:

Corollary 11.7.5. *There exists a polynomial time algorithm that, given a bidirected graph B and signed vertices (x, α) and (y, β) of B with $x \neq y$, either finds an (x, α)–(y, β) trail in B or determines that no such trail exists.*

Proof. Choose an orientation ν of B such that the startvertex of $\nu(e)$ is x for all $e \in E(B)$ incident to x and such that the endvertex of $\nu(f)$ is y for all $f \in E(B)$ incident to y. Let L be the line graph of B with respect to ν. Then the result follows by applying Lemma 11.7.3 to every pair $(e, +)$ and $(f, -)$ of signed vertices of L such that in B the edge e has sign α at x and the edge f has sign β at y; the choice of ν here implies that any such $(e, +)$–$(f, -)$ path in L indeed translates to an x–y trail in B starting in x with sign α and ending in y with sign β. □

As our polynomial time algorithm for Theorem 11.7.1 follows the proof of Theorem 11.5.1, we will also need to compute the appendage of a path in polynomial time. For this, we first prove a lemma which reduces this problem to finding 'ear trails' in polynomial time which are defined as follows:

Let P be a path in a bidirected graph B, and let v be the startvertex of P. Let A be a (P, v)-admissible set. The set $\mathrm{Bon}(P, v, A)$ is then defined to contain $(v, +)$ and $(v, -)$ as well as all signed vertices (w, α) of B for which there exists an oriented edge $\vec{e} \in \overleftrightarrow{A} \cup \vec{E}(P)$ that points at w with sign $-\alpha$. Note that for each $(w, \alpha) \in \mathrm{Bon}(P, v, A)$, there is a v–$(w, -\alpha)$ trail in $A \cup E(P)$. A nontrivial trail T in B is a (P, v, A)-*quasi-ear trail* if both

the signed start- and endvertex of T are contained in $\mathrm{Bon}(P, v, A)$ and the first and last edge of T are not in $\overleftrightarrow{A} \cup \overleftrightarrow{E}(P)$. If this latter condition holds additionally for all edges of T, that is $E(T) \cap (A \cup E(P)) = \emptyset$, then T is an (P, v, A)-*ear trail*.

We observe that for every (P, v, A)-ear trail T in B, the set $A \cup E(T)$ is (P, v)-admissible. Indeed, the signed endvertex (w, α) of T is in $\mathrm{Bon}(P, v, A)$ and thus there exists a v–$(w, -\alpha)$ trail T_w in $A \cup E(P)$. So for every $\vec{e} \in \vec{E}(T)$ the concatenation of $\vec{e}T$ and T_w^- is a $\vec{e}$–v trail in $A \cup E(P) \cup E(T)$. Similarly, there exists an $\vec{e}$–v trail in $A \cup E(P) \cup E(T)$ for every $\vec{e} \in \vec{E}(T^-)$.

Lemma 11.7.6. *Let P be a path in a bidirected graph B, and let v be the startvertex of P. For every non-maximal (P, v)-admissible set $A \subsetneq A(P, v)$, there is either an edge $e \in E(P) \setminus A$ such that $A \cup \{e\}$ is (P, v)-admissible or there exists an (P, v, A)-ear trail.*

Proof. Assume that there is no edge $e \in E(P) \setminus A$ such that $A \cup \{e\}$ is (P, v)-admissible. This implies the following observation which we will use repeatedly throughout this proof: Let Q be a v–$\vec{e}$ trail in $A(P, v) \cup E(P)$ for some orientation $\vec{e}$ of e. Then $\vec{e} \notin \vec{E}(P^-) \setminus \overleftrightarrow{A}$ as otherwise $A \cup \{e\}$ is (P, v)-admissible witnessed by $Q\vec{e}$ and $P\overleftarrow{e}$, a contradiction to the assumption.

The proof now proceeds in two steps. We first find a (P, v, A)-quasi-ear trail in B. Secondly, we show that every (P, v, A)-quasi-ear trail contains a (P, v, A)-ear trail as a subtrail.

Since A is a non-maximal (P, v)-admissible set, there exists $e \in A(P, v) \setminus A$. The above observation then implies that no orientation $\vec{e}$ of e can be in $\vec{E}(P^-)$, since $e \in A(P, v)$ implies the existence of a v–$\vec{e}$ trail in $A(P, v) \cup E(P)$; so we have $e \in A(P, v) \setminus (A \cup E(P))$.

Fix an arbitrary orientation $\vec{e}$ of e. Since $e \in A(P, v)$, there is a v–$\vec{e}$ trail T in $A(P, v) \cup E(P)$, and we let $\vec{f}$ be the first edge on T not contained in $\overleftrightarrow{A} \cup \vec{E}(P)$; note that such an edge $\vec{f}$ exists since $\vec{e}$ is a suitable choice. The above observation then implies $\vec{f} \notin \vec{E}(P^-)$, and we hence have $f \in A(P, v) \setminus (A \cup E(P))$. Thus, there exists a v–$\overleftarrow{f}$ trail T' in $A(P, v) \cup E(P)$ and, as above, we let $\vec{f}'$ be the first edge on T' not contained in $\overleftrightarrow{A} \cup \vec{E}(P)$. Again, we have $\vec{f}' \in A(P, v) \setminus (A \cup E(P))$ by the above observation. By this construction, $\vec{f}'T'\overleftarrow{f}$ is a (P, v, A)-quasi-ear trail in B.

Now let T be a (P, v, A)-quasi-ear trail in B. By definition, T is a (P, v, A)-ear trail if and only if $E(T) \cap (A \cup E(P)) = \emptyset$. So suppose that T is not itself a (P, v, A)-ear trail, and let $\vec{f}T\vec{f}'$ be a maximal subtrail of T in $A \cup E(P)$. Since the first and the last edge of T are not in $\overleftrightarrow{A} \cup \overleftrightarrow{E}(P)$, there exist an edge $\vec{g}$ preceding $\vec{f}$ on T and an edge $\vec{g}'$ succeeding $\vec{f}'$ on T. We now claim that at least one of $T\vec{g}$ and $\vec{g}'T$ is again a (P, v, A)-quasi-ear trail. This then concludes the proof: Since both $T\vec{g}$ and $T\vec{g}'$ contain strictly fewer edges in $A \cup E(P)$, we can recursively apply the above step until we obtain a (P, v, A)–ear trail.

Both $\vec{g}$ and $\vec{g}'$ are not in $\overleftrightarrow{A} \cup \overleftrightarrow{E}(P)$ by construction. So by the definition of $\mathrm{Bon}(P, v, A)$, the claim holds if $\vec{f} \in \overleftrightarrow{A} \cup \vec{E}(P^-)$ or $\vec{f}' \in \overleftrightarrow{A} \cup \vec{E}(P)$. Thus, we can assume that $\vec{f} \in \vec{E}(P) \setminus \overleftrightarrow{A}$ and $\vec{f}' \in \vec{E}(P^-) \setminus \overleftrightarrow{A}$ since $f, f' \in A \cup E(P)$. Let $\vec{h}$ be the first edge of $\vec{E}(\vec{f}T\vec{f}')$ along P. If we have $\vec{h} \in \vec{E}(T)$, then $P\vec{h}T\vec{f}'$ is a trail. This contradicts our above observation since $\vec{f}' \in \vec{E}(P^-) \setminus \overleftrightarrow{A}$. Otherwise, $\vec{h} \in \vec{E}(T^-)$ and then $P\vec{h}T^-\overleftarrow{f}$ is a trail. This contradicts our above observation since $\overleftarrow{f} \in \vec{E}(P^-) \setminus \overleftrightarrow{A}$. □

Corollary 11.7.7. *There exists a polynomial time algorithm that, given a path P with startvertex v in a bidirected graph B, computes its appendage $A(P, v)$.*

Proof. The algorithm starts with the (P, v)-admissible set $A = \emptyset$ and extends A recursively to (P, v)-admissible sets until $A = A(P, v)$. Let A be any (P, v)-admissible set.

The algorithm first checks whether there exists an edge $e \in E(P) \setminus A$ such that $A \cup \{e\}$ is again (P, v)-admissible. Such an edge e exists if and only if there exists a v–$\vec{e}$ trail in $A \cup E(P)$ where $\vec{e} \in \overrightarrow{E}(P^-)$. This can be checked in polynomial time by Corollary 11.7.5. If such an edge e exists, then the algorithm restarts the recursion with $A \cup \{e\}$.

If there is no such edge, then the algorithm looks for a (P, v, A)-ear trail T in B. This can again be done in polynomial time by Corollary 11.7.5 since there exists such a (P, v, A)-ear trail if and only if there is a nontrivial trail in $B - (A \cup E(P))$ that starts and ends in $\mathrm{Bon}(P, v, A)$. If such T exists, then the algorithm restarts the recursion with $A \cup E(T)$ which is again (P, v)-admissible by the definition of (P, v, A)-ear trail.

If there is also no (P, v, A)-ear trail, then Lemma 11.7.6 implies $A = A(P, v)$, and the algorithm terminates. Since the number of recursion steps is bounded by $|E(B)|$, this algorithm runs in polynomial time. □

With all these polynomial time algorithms at hand, we are now ready to prove Theorem 11.7.1:

Proof of Theorem 11.7.1. Our algorithm follows the construction in the proof of Theorem 11.5.1 which thus yields the correctness of the algorithm. Each construction step of the recursive procedure in the proof of Theorem 11.5.1 can be done in polynomial time by Lemma 11.7.3, Corollary 11.7.5 and Corollary 11.7.7. The number of recursion steps is at most the sum of the lengths of the paths P_i and hence bounded by $|E(B)|$. Altogether, this yields a polynomial running time. □

CHAPTER 12

Hitting cycles through prescribed vertices or edges

This chapter is based on [BGGJR24] and joint work with Nathan Bowler, Ebrahim Ghorbani, Florian Gut and Florian Reich.

We prove that for every set S of vertices of a directed graph D, the maximum number of vertices in S contained in a collection of vertex-disjoint cycles in D is at least the minimum size of a set of vertices that hits all cycles containing a vertex of S. As a consequence, the directed tree-width of a directed graph is linearly bounded in its cycle-width, which improves the previously known quadratic upper bound.

We further show that the corresponding statement in bidirected graphs is true and that its edge-variant holds in both undirected and directed graphs, but fails in bidirected graphs. The vertex-version in undirected graphs remains an open problem.

12.1 Introduction

How many vertices of a graph G does it take to hit all cycles in G? The *Feedback Vertex Set Problem* which looks for a minimal such set of vertices of G is among Karp's classical NP-complete problems [K72]. The celebrated Erdős-Pósa Theorem asserts a qualitative duality between feedback vertex sets and collections of vertex-disjoint cycles:

Theorem 12.1.1 ([EP65])**.** *There exists $c \in \mathbb{N}$ such that, for each $k \in \mathbb{N}$, every undirected graph G contains either k vertex-disjoint cycles or a set X of at most $c\,k \log k$ vertices that hits every cycle of G.*

In the context of directed graphs the duality between feedback vertex set and vertex-disjoint cycles is known as Younger's conjecture and was proven by Reed, Robertson, Seymour and Thomas:

Theorem 12.1.2 ([RRST96])**.** *For each $k \in \mathbb{N}$ there exists an integer $d \in \mathbb{N}$ such that every directed graph D contains either k vertex-disjoint cycles or a set X of at most d vertices that hits every cycle of D.*

The duality between feedback edge sets and edge-disjoint cycles in undirected and directed graphs holds true in the same way (by applying a simple auxiliary construction [PW12] or the directed line graph [HU18], respectively).

More generally, one can consider *S-cycles*, that is, cycles hitting a prescribed set S of vertices. Wollan and Pontecorvi extended Theorem 12.1.1 to S-cycles:

Theorem 12.1.3 ([PW12]). *There exists $c \in \mathbb{N}$ such that, for each $k \in \mathbb{N}$, undirected graph G and set $S \subseteq V(G)$, the graph G contains either k vertex-disjoint S-cycles or a set X of at most $c\,k \log k$ vertices hitting every S-cycle in G.*

Kakimura and Kawarabayashi extended Theorem 12.1.2 fractionally to S-cycles [KK12], and Kawarabayashi, Král', Krčál and Kreutzer later did so half-integrally [KKKK13]. Both these results as well as Theorem 12.1.3 transfer to edge-disjoint cycles through prescribed sets of edges (again by applying a simple auxiliary construction [PW12] or the directed line graph, respectively).

Given a set S of vertices of an undirected graph, every set of k vertex-disjoint S-cycles clearly witnesses that a minimum feedback vertex set for the graph's S-cycles has size at least k. The converse of this does not hold, as Theorem 12.1.3 is sharp: there exist graphs and vertex sets S with at most k vertex-disjoint S-cycles, where every feedback vertex set for the S-cycles has size at least $\Omega(k \log k)$. But instead we can ask for a set of vertex-disjoint S-cycles which meets S as often as the minimum size of a feedback vertex set for the S-cycles.

To state this formally, we write, for a collection $\mathcal{C}$ of cycles in an (un)directed graph, $V(\mathcal{C})$ for the set of all vertices that are contained in some cycle in $\mathcal{C}$ and $E(\mathcal{C})$ for the analogous set of edges.

Question 12.1.4. *Let G be an undirected graph and let $S \subseteq V(G)$. Is it true that*

$$\max\{|S \cap V(\mathcal{C})| \colon \mathcal{C} \text{ is a collection of vertex-disjoint cycles of } G\}$$
$$\geqslant \min\{|X| \colon X \subseteq V(G) \text{ hits all } S\text{-cycles in } G\}?$$

The same question arises for vertex-disjoint cycles in directed and bidirected graphs as well as for edge-disjoint F-cycles meeting a prescribed set of edges F.

In this chapter, we study these various versions of Question 12.1.4 and obtain the following results:

	undirected graph	**directed graph**	**bidirected graph**
vertex-version	*open*	*true*, see Theorem 12.2.7	*true*, see Theorem 12.2.3
edge-version	*true*, see Theorem 12.2.1	*true*, see Theorem 12.2.2	*false*, see Proposition 12.2.8

To our surprise, we have not been able to deduce the undirected statement from its directed analogue. The naive approach suggests that we should replace every undirected edge by two parallel directed edges with distinct directions, but this creates many new cycles (of length 2) and as such gives no direct correspondence between the cycles in the undirected graph and the cycles in the corresponding directed graph. Similar methods which replace the two parallel directed edges by a more involved 'gadget' seem to fail for the same reason. Therefore, Question 12.1.4 remains an open conjecture.

The proofs of the various versions of Question 12.1.4 mostly rely on algebraic methods. For the edge-version in undirected graphs, Theorem 12.2.1, the proof uses the cycle space over $\mathbb{Z}_2$ as a key tool. The proof of the edge-version in directed graphs, Theorem 12.2.2, is based on the strong duality theorem for linear programming. The vertex-version in directed graphs, Theorem 12.2.7, follows from the respective statement in bidirected

graphs, Theorem 12.2.3, whose proof also relies on the strong duality theorem for linear programming.

As an application of the vertex-version in directed graphs, Theorem 12.2.7, we improve a bound between two different width-measures on directed graphs, directed tree-width and cycle-width. *Directed tree-width* is the well-established directed analogue of tree-width in undirected graphs [JRST01], exhibiting similar features such as a directed grid theorem using butterfly minors [KK15] and a respective cops & robbers-game [JRST01]. The structure of directed graphs is known to be closely connected to the structure of bipartite graphs with perfect matchings; see e.g. [W22] and especially its Section 3.2 for an overview. One of the central open questions on the structure of graphs with perfect matchings is Norine's matching grid conjecture [N05], which asserts an analogue of the (directed) grid theorem in terms of perfect matching width and matching minors. The notion of *cycle-width* was introduced by Hatzel, Rabinovich and Wiederrecht [HRW19] to bridge the gap between directed tree-width and perfect matching width. Roughly speaking, they showed that cycle-width is equivalent to both directed tree-width and perfect matching width, which allowed them to apply the directed grid theorem to resolve Norine's matching grid conjecture for bipartite graphs.

Cycle-width is a branch decomposition-type width-measure of directed graphs where the order of a cut measures how often a family of vertex-disjoint cycles can meet the edges in the cut (see Definition 12.3.3 for the formal definition). Together with the cops & robbers-game corresponding to directed tree-width [JRST01], Theorem 12.2.7 allows us to show the following upper bound on directed tree-width in terms of cycle-width:

Corollary 12.1.5. *For every directed graph D, we have* $\mathrm{dtw}(D) \leqslant 9\,\mathrm{cycw}(D) - 2$.

This improves the previously known quadratic upper bound for directed graphs by Giannopoulou, Kreutzer and Wiederrecht [GKW24] and, together with their lower bound, implies that directed tree-width and cycle-width are linearly bounded in each other. Moreover, our proof of Corollary 12.1.5 is significantly simpler than their technically involved proof. We note that their proof contains a variant of Question 12.1.4 in directed graphs with weaker bounds, which inspired us to the present work.

Bidirected graphs can be obtained from undirected graphs by assigning to each endvertex of an edge one of two signs, $+$ or a $-$; directed graphs can then be viewed as those bidirected graphs in which each edge has sign $+$ at one endvertex and sign $-$ at the other (see Section 12.2.3 for the formal definition). While directed graphs characterise bipartite graphs with perfect matchings, bidirected graphs play the same role for general graphs with perfect matchings [W22]. This is why, in his doctoral thesis [W22], Wiederrecht suggested the investigation of structural properties of bidirected graphs as a route towards the full resolution of Norine's matching grid conjecture [N05].

Following this route, we also study Question 12.1.4 in bidirected graphs, and we show with Theorem 12.2.3 that the vertex-version in bidirected graphs is true. As for directed graphs, Theorem 12.2.3 implies a variant of Corollary 12.1.5 for bidirected graphs, i.e. bidirected tree-width is linearly bounded by cycle-width (Corollary 12.3.2). In contrast, the edge-version in bidirected graphs does not hold (Proposition 12.2.8). Our counterexample

builds on the fact that Menger's theorem does not transfer to bidirected graphs without additional assumptions [BGGJR23].

This chapter is structured as follows. We first state and prove our results on the various versions of Question 12.1.4 in Section 12.2. In Section 12.3 we then discuss the concepts of (bi)directed tree-width and cycle-width and prove Corollary 12.1.5 and its bidirected analogue, Corollary 12.3.2.

For all graph-theoretic terms and notations, we follow [D25].

12.2 Proofs of the main results of this chapter

In this section, we state and prove the main results of this chapter.

12.2.1 The edge-version in undirected graphs

Theorem 12.2.1. *Let G be an undirected graph, and let $F \subseteq E(G)$ be a set of edges. Then*

$$\begin{aligned}&\max\{|F \cap E(\mathcal{C})| \colon \mathcal{C} \text{ is a collection of edge-disjoint cycles of } G\}\\ &\geqslant \min\{|Y| \colon Y \subseteq E(G) \text{ hits all } F\text{-cycles in } G\}.\end{aligned}$$

Proof. Let $\mathcal{A}$ be the set of all F-cycles in G, and let M be the incidence matrix of $E(\mathcal{A})$ versus $\mathcal{A}$: The rows and columns of M correspond to the edges in $E(\mathcal{A})$ and the cycles in $\mathcal{A}$, respectively. For $e \in E(\mathcal{A})$ and $C \in \mathcal{A}$, the entry at position (e, C) then is 1 if $e \in C$ and 0 if $e \notin C$.

The column space of M over the field $\mathbb{Z}_2$, denoted as $\mathrm{Col}_{\mathbb{Z}_2}(M)$, is a subspace[1] of the cycle space of the induced subgraph of G on $E(\mathcal{A})$. This implies that every vector in $\mathrm{Col}_{\mathbb{Z}_2}(M)$ represents a collection of edge-disjoint cycles. Indeed, the mod-2 sum of some columns of M, corresponding for example to cycles $C_1, \dots, C_\ell \in \mathcal{A}$, gives the vector representing the symmetric difference $E(C_1) \triangle \cdots \triangle E(C_\ell)$, which induces a collection of edge-disjoint cycles in G.

Now let t be the value on the right-hand side of the inequality in the statement. By the above, it suffices to find a vector in $\mathrm{Col}_{\mathbb{Z}_2}(M)$ whose restriction to F has at least t many 1's, which is what we prove below.

For every $E' \subseteq E(\mathcal{A})$ let us denote by $M_{E'}$ the submatrix of M consisting of the rows corresponding to edges in E'. We claim that the rank of M_F over $\mathbb{Z}_2$ is at least t. Otherwise, there exists $F' \subseteq F$ of size $< t$ such that every row of M_F is the mod-2 sum of the rows of $M_{F'}$. This implies that all cycles in $\mathcal{A}$ meet F': Every column of M_F is non-zero by the definition of M, and since the rows of $M_{F'}$ generate every other row of M_F, every column of M_F must have at least one entry equal to 1 in $M_{F'}$-part. In other words, F' hits all cycles in $\mathcal{A}$ and has size $< t$, contradicting the minimality of t.

Since M_F has rank at least t over $\mathbb{Z}_2$, we may assume that the matrix M has the

[1] Indeed, $\mathrm{Col}_{\mathbb{Z}_2}(M)$ can be proved to be equal to the cycle space of the induced subgraph of G on $E(\mathcal{A})$. However, we do not need this property in our argument.

following form in which the $t \times t$ submatrix $N_{t\times t}$ is non-singular over $\mathbb{Z}_2$:

$$M = \begin{array}{|c|c|}\hline N_{t\times t} & \multirow{2}{*}{M_F} \\ \cline{1-1} & \\ \hline & \\ \hline \end{array}.$$

But since $N_{t\times t}$ is non-singular, its columns can generate every vector in $\mathbb{Z}_2^t$ and in particular the all-1 vector of length t. Thus, there is a vector in $\mathrm{Col}_{\mathbb{Z}_2}(M)$ whose first t entries are 1, as desired. □

12.2.2 The edge-version in directed graphs

Theorem 12.2.2. *Let D be a directed graph, and let $F \subseteq E(D)$ be a set of edges. Then*

$$\begin{aligned}&\max\{|F \cap E(\mathcal{C})|\colon \mathcal{C} \text{ is a collection of edge-disjoint cycles of } D\}\\ &\geqslant \min\{|Y|\colon Y \subseteq E(D) \text{ hits all } F\text{-cycles in } D\}.\end{aligned} \tag{12.1}$$

Proof. The proof relies on the strong duality theorem of linear programming. To this end, we construct a primal linear program whose optimal value equals the left-hand side of (12.1), and the optimal value of the dual is an upper bound for the right-hand side of (12.1).

Let n be the number of vertices of D and let m be the number of edges of D. Further, let $N_{n\times m}$ be the incidence matrix of D, and set

$$A := \begin{bmatrix} I_m \\ N \end{bmatrix},$$

where I_m is the identity matrix of size m. Let $\mathbf{f}$ be the characteristic vector of F as a subset of $E(D)$, and let

$$\mathbf{b} := \begin{bmatrix} \mathbf{1}_m \\ \mathbf{0}_n \end{bmatrix},$$

where $\mathbf{1}_m$ is the all-1 vector of length m. Consider the primal linear program

$$\begin{aligned}&\text{maximise } \mathbf{f}^\intercal \mathbf{x}\\ &\text{subject to } A\mathbf{x} \leqslant \mathbf{b}\\ &\qquad\qquad \mathbf{x} \geqslant 0.\end{aligned}$$

Recall that an integer-valued matrix is totally unimodular if every non-singular square submatrix has determinant $+1$ or -1. The incidence matrix of every directed graph, and thus N, is known to be totally unimodular [B93, Proposition 5.3]. It turns out that A is also totally unimodular. To see this, consider an arbitrary square submatrix R of A. If R is itself a submatrix of N or of I_m, or if R contains a row or a column consisting entirely of zeros, then $\det R \in \{0, \pm 1\}$ and there is nothing to prove. Otherwise, by expanding the determinant of R along the rows corresponding to the I_m block, we obtain $\det R = \pm \det R'$, where R' is a square submatrix of N. Since N is totally unimodular, it follows that $\det R \in \{0, \pm 1\}$, and hence A is totally unimodular. Therefore, both the primal linear program and its dual have integral optimal solutions (see e.g. [S98, Corollary 7.1g]).

Let $\mathbf{x}^*$ be an integral optimal solution of the primal linear program. The constraints demand $\mathbf{x}^*$ to be a $\{0,1\}$-vector in the null-space of N. Consequently, the set of edges $\{e \in E(D) : \mathbf{x}^*(e) = 1\}$ induces a directed subgraph of D in which, at every vertex, the number of incoming edges equals the number of outgoing edges. Therefore, $\mathbf{x}^*$ represents the set of edges of some edge-disjoint cycles in D, where its optimality implies that it has the maximum intersection with F. Thus, $\mathbf{f}^\intercal\mathbf{x}^*$ gives the left-hand side of (12.1).

Now we consider the dual program

$$\begin{aligned} &\text{minimise } \mathbf{b}^\intercal\mathbf{y} \\ &\text{subject to } A^\intercal\mathbf{y} \geqslant \mathbf{f} \\ &\qquad\qquad \mathbf{y} \geqslant 0. \end{aligned}$$

Let

$$\mathbf{y}^* := \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{bmatrix}$$

be an integral optimal solution of the dual linear program, where $\mathbf{y}_1$ has length m and $\mathbf{y}_2$ has length n. Let Y be the set of edges in D whose corresponding components in $\mathbf{y}_1$ are non-zero. As $\mathbf{y}_1$ is a non-negative integral vector, we have $|Y| \leqslant \mathbf{1}^\intercal\mathbf{y}_1 = \mathbf{b}^\intercal\mathbf{y}^*$. By the strong duality theorem of linear programming (see e.g. [S98, Corollary 19.2b]), $\mathbf{b}^\intercal\mathbf{y}^* = \mathbf{f}^\intercal\mathbf{x}^*$. So we will be done by showing that Y hits all F-cycles.

Let C be a cycle with the characteristic vector $\mathbf{c}$ (as a subset of $E(D)$). We can write $A^\intercal\mathbf{y}^* \geqslant \mathbf{f}$ as $\mathbf{y}_1 + N^\intercal\mathbf{y}_2 \geqslant \mathbf{f}$. Multiplying both sides by $\mathbf{c}^\intercal$ from the left, and observing that $\mathbf{c}^\intercal N^\intercal = \mathbf{0}$, we get $\mathbf{c}^\intercal\mathbf{y}_1 \geqslant \mathbf{c}^\intercal\mathbf{f}$. If C hits F, then $\mathbf{c}^\intercal\mathbf{f} = |E(C) \cap F| > 0$, implying $\mathbf{c}^\intercal\mathbf{y}_1 > 0$, which means that C hits Y as well. □

12.2.3 The vertex-version in bidirected graphs

Let us briefly recall the relevant definitions around bidirected graphs; for a complete formal introduction, we refer the reader to [BGGJR23].

A *bidirected graph* $B = (G, \sigma)$ consists of a graph G, which is undirected and may contain parallel edges but no loops, and a *signing* σ which assigns to each half-edge (u, e) of G, where $e \in E(G)$ and $u \in e$, one of the signs $+$ and $-$; we say that e *has sign* $\sigma(u, e)$ *at* u. All terms such as vertices, edges, adjacent, incident, edge- and vertex-disjoint transfer directly from G to B. A *path* in B is a path in G such that for every internal vertex v of the path, the two incident edges have distinct signs at v; in other words, the sign switches at each internal vertex of the path. *Cycles* in B are then the analogous sign-switching version of cycles in G.

Theorem 12.2.3. *Let B be a bidirected graph and let $S \subseteq V(B)$ be a set of vertices. Then*

$$\begin{aligned} &\max\{|V(\mathcal{C}) \cap S| : \mathcal{C} \text{ a collection of vertex-disjoint cycles}\} \\ &\geqslant \min\{|X| : X \subseteq V(D) \text{ hits all } S\text{-cycles}\}. \end{aligned}$$

The proof of Theorem 12.2.3 applies linear programming in a similar way to the proof of Theorem 12.2.2, i.e. solutions to the primal program correspond to edge-disjoint cycles. For this we transfer the vertex-problem into an edge-problem: Let B' be the bidirected

graph obtained from B by replacing each vertex $v \in V(B)$ with two vertices v^+, v^- and making v^+ incident with all edges incident with v with a plus sign and making v^- incident with all edges incident with v with a minus sign. Furthermore, connect v^- and v^+ by an edge with minus sign at v^+ and plus sign at v^-.

We observe that:

(i) there is a one-to-one correspondence between the cycles of B and those of B';

(ii) two cycles in B are vertex-disjoint if and only if the corresponding cycles in B' are vertex-disjoint;

(iii) every two edge-disjoint cycles in B' are also vertex-disjoint.

Although the edge-version in bidirected graphs is not true in general (see Section 12.2.5), for such a bidirected graph B' as constructed above and the set of edges induced by S it is equivalent to the vertex-version in bidirected graphs and we will show it to hold.

As a further ingredient to the proof of Theorem 12.2.3, we need the following concepts around incidence matrices of bidirected graphs. The *incidence matrix* of a bidirected graph (G, σ) is defined based on the signs of the 'half-edges', where the sign of the entries $+1$ or -1 corresponding to a vertex v and an edge e is equal to $\sigma(v, e)$.

Remark 12.2.4. For the bidirected graph B', every $\{0, 1\}$-vector $\mathbf{x}$ in the null-space of its incidence matrix corresponds to the edge set of some union of vertex-disjoint cycles in B' (this is because, in the subgraph induced by $\{e \in E(B') : \mathbf{x}^*(e) = 1\}$, the edge between v^- and v^+ can be used at most once, which implies that each vertex in this subgraph has exactly two incident edges). This property fails for general bidirected graphs (see Figure 12.1 for an example), which explains why the following proof will not work for the general edge-version problem.

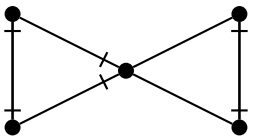

Figure 12.1: A bidirected graph for which the columns in its incidence matrix sum up to zero because each vertex has an equal number of positive and negative incident edges, yet it has no bidirected cycle.

While the total unimodularity of incidence matrices of directed graphs relies on the characteristic property that each column has exactly one $+1$ and one -1 entry, this characteristic does not necessarily hold for bidirected graphs, whose incidence matrices may thus not be totally unimodular. Despite this, the following more general notions allow us to extend the integrality of optimal solutions of the linear programs considered for directed graphs to bidirected graphs.

A rational matrix M is called *k-regular*, for a positive integer k, if for every non-singular submatrix R of M, kR^{-1} is integral. If M is k-regular, then the matrix $\frac{1}{k}M$ is 1-regular. Indeed, every non-singular submatrix of $\frac{1}{k}M$ is of the form $\frac{1}{k}R$, where R is a non-singular submatrix of M, and its inverse is given by $\left(\frac{1}{k}R\right)^{-1} = kR^{-1}$ which is integral.

Lemma 12.2.5 ([AK04, Theorem 23])**.** *The incidence matrix of a bidirected graph is* 2*-regular.*

Further, a polyhedron $P = \{\mathbf{x} : \mathbf{x} \geqslant \mathbf{0}, M\mathbf{x} \leqslant \mathbf{a}\}$ is called *integral* if $\max\{\mathbf{c}^\intercal\mathbf{x} : \mathbf{x} \in P\}$ is attained by an integral vector for each $\mathbf{c}$ for which the maximum is finite.

Lemma 12.2.6 ([AK04, Theorem 16])**.** *Let $M_{n\times m}$ be a rational matrix. Then M is* 1*-regular if and only if for every integral vector $\mathbf{a}$ of length n the polyhedron $\{\mathbf{x} : \mathbf{x} \geqslant \mathbf{0}, M\mathbf{x} \leqslant \mathbf{a}\}$ is integral.*

Proof of Theorem 12.2.3. Let n be the number of vertices of B' and let m be the number of edges of B'. Further, let $F \subseteq E(B')$ be the set of all edges constructed from vertices in $S \subseteq V(B)$. It suffices to establish the following:

$$\begin{aligned}&\max\{|\mathcal{C} \cap F| : \mathcal{C} \text{ a set of edge-disjoint cycles in } B'\}\\ &\geqslant \min\{|Y| : Y \subseteq E(B') \text{ hits all } F\text{-cycles in } B'\}.\end{aligned} \tag{12.2}$$

Let $\mathbf{f}$ be the characteristic vector of F as a subset of $E(B')$, and let

$$\mathbf{b} := \begin{bmatrix}\mathbf{1}_m\\ \mathbf{0}_n\\ \mathbf{0}_n\end{bmatrix}.$$

Also let M be the incidence matrix of B', and

$$A := \begin{bmatrix}I_m\\ \frac{1}{2}M\\ -\frac{1}{2}M\end{bmatrix}.$$

Consider the primal linear program

$$\begin{aligned}&\text{maximise } \mathbf{f}^\intercal\mathbf{x}\\ &\text{subject to } A\mathbf{x} \leqslant \mathbf{b}\\ &\qquad\qquad \mathbf{x} \geqslant 0.\end{aligned}$$

By Lemma 12.2.5, M is 2-regular and hence $\frac{1}{2}M$ is 1-regular. It follows that $\begin{bmatrix}\frac{1}{2}M\\ -\frac{1}{2}M\end{bmatrix}$ is also 1-regular because every non-singular submatrix of it can be viewed as a submatrix of $\frac{1}{2}M$ with some rows possibly negated. Stacking an identity matrix on a 1-regular matrix again yields a 1-regular matrix (see [AK04, Lemma 7]), so A is 1-regular as well.

According to Lemma 12.2.6, the polyhedron $P := \{\mathbf{x} : \mathbf{x} \geqslant \mathbf{0}, A\mathbf{x} \leqslant \mathbf{b}\}$ is integral. Moreover, $\max\{\mathbf{f}^\intercal\mathbf{x} : \mathbf{x} \in P\}$ is finite, since $A\mathbf{x} \leqslant b$ and $\mathbf{x} \geqslant 0$ require all entries of $\mathbf{x}$ to be contained in the interval $[0, 1]$. Thus, the primal linear program has an integral optimal solution $\mathbf{x}^*$. The constraints demand $\mathbf{x}^*$ to be a $\{0, 1\}$-vector in the null-space of M. Therefore, by Remark 12.2.4, $\mathbf{x}^*$ represents the set of edges of some union of edge-disjoint cycles in B', where its optimality implies that it has the maximum intersection with F. Thus, $\mathbf{f}^\intercal\mathbf{x}^*$ gives the left-hand side of (12.2).

Now we consider the dual program

$$\begin{aligned} &\text{minimise } \mathbf{b}^\intercal \mathbf{y} \\ &\text{subject to } A^\intercal \mathbf{y} \geqslant \mathbf{f} \\ &\mathbf{y} \geqslant 0, \end{aligned}$$

and show that it has an integral optimal solution. By the strong duality theorem of linear programming (see e.g. [S98, Corollary 19.2b]), the dual program has an optimal solution $\mathbf{y}*$ with $\mathbf{b}^\intercal \mathbf{y}^* = \mathbf{f}^\intercal \mathbf{x}^*$. The matrix $-A^\intercal$ is 1-regular since A is 1-regular. Thus, by Lemma 12.2.6, the polyhedron $\{\mathbf{x} : \mathbf{x} \geqslant \mathbf{0}, -A^\intercal \mathbf{x} \leqslant -\mathbf{f}\} = \{\mathbf{y} : \mathbf{y} \geqslant \mathbf{0}, A^\intercal \mathbf{y} \leqslant \mathbf{f}\}$ is integral. Since maximising $-\mathbf{b}^\intercal \mathbf{y}$ is equal to minimising $\mathbf{b}^\intercal \mathbf{y}$, we may hence assume $\mathbf{y}^*$ to be an integral optimal solution

$$\mathbf{y}^* = \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_3 \end{bmatrix}$$

for the dual program, where $\mathbf{y}_1$ has length m and both $\mathbf{y}_2$ and $\mathbf{y}_3$ have length n. Let Y be the set of edges in B' whose corresponding components in $\mathbf{y}_1$ are non-zero. As $\mathbf{y}_1$ is a non-negative integral vector, we have $|Y| \leqslant \mathbf{1}^\intercal \mathbf{y}_1 = \mathbf{b}^\intercal \mathbf{y}^*$. Since $\mathbf{b}^\intercal \mathbf{y}^* = \mathbf{f}^\intercal \mathbf{x}^*$, we will thus be done by showing that Y hits all F-cycles.

Let C be an arbitrary cycle in B' and write $\mathbf{c}$ for the characteristic vector as a subset of $E(B')$. We can write $A^\intercal \mathbf{y}^* \geqslant \mathbf{f}$ as $\mathbf{y}_1 + \frac{1}{2} M^\intercal (\mathbf{y}_2 - \mathbf{y}_3) \geqslant \mathbf{f}$. Multiplying both sides by $\mathbf{c}^\intercal$ from the left, and observing that $\mathbf{c}^\intercal M^\intercal = \mathbf{0}$, we get $c^\intercal \mathbf{y}_1 \geqslant \mathbf{c}^\intercal \mathbf{f}$. If C hits F, then $\mathbf{c}^\intercal \mathbf{f} = |E(C) \cap F| > 0$, implying $\mathbf{c}^\intercal \mathbf{y}_1 > 0$, which means C hits Y as well.

So we have proved that $|Y|$ is at most the left-hand side of (12.2), and at least the right-hand side of (12.2), completing the proof. □

12.2.4 The vertex-version in directed graphs

Every directed graph D forms a bidirected graph in the following way: Let G be the underlying undirected graph of D and let σ be the map sending a half edge (u, e) to $+$ if u is the head of e and to $-$ if u is the tail of e. Then (G, σ) is a bidirected graph whose cycles correspond one-to-one to the cycles of D.

We can deduce from Theorem 12.2.3:

Theorem 12.2.7. *Let D be a directed graph, and let $S \subseteq V(D)$ be a set of vertices. Then*

$$\begin{aligned} &\max\{|S \cap V(\mathcal{C})| : \mathcal{C} \text{ a collection of vertex-disjoint cycles in } D\} \\ &\geqslant \min\{|X| : X \subseteq V(D) \text{ hits all } S\text{-cycles in } D\}. \end{aligned}$$

12.2.5 A counterexample to the edge-version in bidirected graphs

The following proposition now asserts that the edge-version of Question 12.1.4 does not hold in bidirected graphs:

Proposition 12.2.8. *For every $k \in \mathbb{N}$ there is a bidirected graph B_k and a set F_k of edges of B_k such that*

(1) for every collection $\mathcal{C}$ of edge-disjoint cycles in B_k, we have $|F_k \cap E(\mathcal{C})| \leqslant 1$, and

(2) for every edge set X_k of size at most k, $B_k - X_k$ contains a cycle meeting F_k.

For the proof of Proposition 12.2.8, we use the following result:

Theorem 12.2.9 ([BGGJR23, Theorem 3.3])**.** *For each number $k \in \mathbb{N}$, there exists a bidirected graph A_k and distinct vertices x and y, such that there are no two edge-disjoint x–y paths and for each subset $S \subseteq E(A_k)$ of size k there exists an x–y path in $A_k - S$.*

Proof of Proposition 12.2.8. Let A_k, x and y be as in Theorem 12.2.9. We may assume without loss of generality that all edges incident to x in A_k have sign $+$ at x and that all edges incident to y in A_k have sign $-$ at y. Now let B_k be the bidirected graph obtained from A_k by identifying x and y, and let F_k be the set of edges with sign $-$ at $x = y$, i.e. the set of edges incident to y in A_k. For an illustration of B_k, refer to Figure 12.2. In this figure, the blue vertex at the top-left corner represents the vertex resulting from the identification of x and y and F_k consists of the dotted edges which were initially incident to y. We claim that B_k and F_k are as desired.

For (1), let C be an arbitrary cycle in B_k containing an edge of F_k (like e.g. the one in red in Figure 12.2). The cycle C contains exactly one edge with sign $-$ at $x = y$ and exactly one edge with sign $+$ at $x = y$; in particular, C contains exactly one edge of F_k. Furthermore, the edge set $E(C)$ induces an x–y path in A_k. Since there are no two edge-disjoint x–y paths in A_k, every collection $\mathcal{C}$ of edge-disjoint cycles has at most one element that hits $x = y$. In particular, $|F_k \cap E(\mathcal{C})| \leqslant 1$ holds for every collection $\mathcal{C}$ of edge-disjoint cycles in B_k.

For (2), let $X_k \subseteq E(B_k) = E(A_k)$ with $|X_k| \leqslant k$ be arbitrary. There is an x–y path P in $A_k - X_k$ by Theorem 12.2.9. By the construction of B_k, the edge set $E(P)$ induces a cycle in $B_k - X_k$, and this cycle contains an edge of F_k. □

12.3 Application: (bi)directed tree-width and cycle-width

Cycle-width was introduced by Hatzel, Rabinovich and Wiederrecht [HRW19] as a tool to make the directed grid theorem of Kawarabayashi and Kreutzer [KK15] accessible in the realm of bipartite graphs with perfect matchings. They used it to resolve Norine's matching grid conjecture [N05] for the case of bipartite graphs. The key features of cycle-width are that it is equivalent to both a respective matching-theoretic width-measure, called perfect matching width, and directed tree-width. Here, we are concerned with this later equivalence between cycle-width and directed tree-width, which we show to be linear.

Corollary 12.3.1. *For every directed graph D, we have* $\mathrm{dtw}(D) \leqslant 9\,\mathrm{cycw}(D) - 2$.

With the same proof, we will also show the corresponding statement for bidirected graphs, namely that bidirected tree-width is linearly bounded by cycle-width. We hope this to be useful towards the full resolution of Norine's matching grid conjecture [N05] along the route suggested by Wiederrecht [W22]. Our statement holds for bidirected graphs $B = (G, \sigma)$ that are *circular*, that is, every edge of B is contained in a cycle and

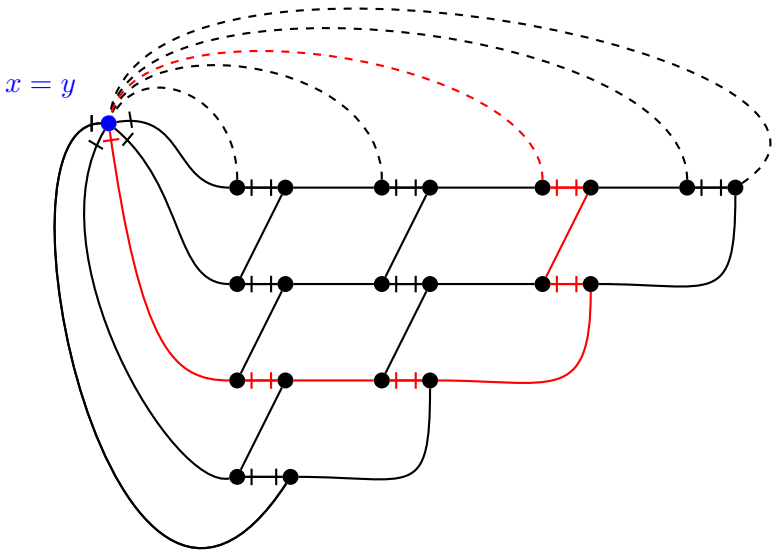

Figure 12.2: A counterexample to the edge-version in bidirected graphs. More precisely, a bidirected graph in which there are no two edge-disjoint cycles containing a dotted edge. For every edge set X of size $\leqslant 2$ there exists a cycle in $B - X$ containing a dotted edge.

G is connected. This, however, is no restriction in the context of strong connectivity in bidirected graphs (see [W22, text before Lemma 9.2.18]).

Corollary 12.3.2. *For every circular bidirected graph B, we have* $\mathrm{btw}(B) \leqslant 9\,\mathrm{cycw}(B)+2$.

Towards the proof of Corollary 12.1.5 and Corollary 12.3.2, we begin with the formal definition of the cycle-width of a (bi)directed graph.

Definition 12.3.3 ([W22, Definition 5.3.11])**.** A *cycle-decomposition* of a (bi)directed graph D is a pair (T, φ) of a cubic tree T and a bijection φ between $V(D)$ and the set $L(T)$ of leaves of T. Each edge t_1t_2 of T then *induces* a cut $E_G(V_1, V_2)$ of D where $V_i = \varphi(L(T_i))$ and T_i is the subtree of $T - t_1t_2$ containing t_i for $i = 1, 2$.

The *cycle porosity* of a cut F of D is defined as

$$\mathrm{cp}(F) := \max_{\substack{\mathcal{C} \text{ collection of} \\ \text{pairwise vertex-disjoint} \\ \text{cycles in } D}} |F \cap E(\mathcal{C})|.$$

The *width* of a cycle-decomposition (T, φ) of D is then defined as the maximum cycle porosity of a cut of D induced by an edge of T. The *cycle-width* $\mathrm{cycw}(D)$ of D is the minimum width of a cycle-decomposition of D.

Our proof of Corollary 12.1.5 and Corollary 12.3.2 proceeds analogously to the one of Giannopoulou, Kreutzer and Wiederrecht in their quadratic upper bound [GKW24, Proposition 3.3]; we include the full formal proof here for completeness. The proof does not use the definition of (bi)directed tree-width itself, but a sufficient condition for bounded (bi)directed tree-width via a suitable 'cops & robbers-game' on (bi)directed

graphs, Theorems 12.3.5 and 12.3.6 below. We therefore use (bi)directed tree-width as a black box here and refer the reader to [W22, Definition 9.2.17] for the definition.

Definition 12.3.4 ([JRST01],[W22, Definition 9.2.11])**.** The *cops & robbers-game* is played between k cops and a robber on a (bi)directed graph D as follows: For the start the cops take an initial position $C_0 \subseteq V(D)$ and the robber chooses her initial position R_0, a strong component of $D - C_0$.

After the i-th round, let $C_i \subseteq V(D)$ be the position of the cops and let R_i be the strong component of $D - C_i$ in which the robber is situated. Since there are k cops, we stipulate that $|C_i| \leqslant k$. The cops then announce their new position $C_{i+1} \subseteq V(D)$ and the robber afterwards chooses as her position a strong component R_{i+1} of $D - C_{i+1}$ such that $R_i \cup R_{i+1} \subseteq R$ for some strong component R of $D - (C_i \cap C_{i+1})$. If no such R_{i+1} exists, then the cops *catch* the robber after $i + 1$ rounds.

If k cops can ensure that they catch the robber after a finite number of rounds independently of the robber's moves, then the cops have a *winning strategy*, and we say that *k cops can catch the robber on B.*

Theorem 12.3.5 ([JRST01])**.** *For every directed graph D and every $k \in \mathbb{N}$, we have* $\mathrm{dtw}(D) \leqslant 3k - 2$ *if k cops can catch the robber on D.*

Theorem 12.3.6 ([W22, Lemma 9.2.13, Theorem 9.2.20])**.** *For every circular bidirected graph B and every $k \in \mathbb{N}$, we have* $\mathrm{btw}(B) \leqslant 3k + 2$ *if k cops can catch the robber on B.*

Proof of Corollary 12.1.5 and Corollary 12.3.2. Before we start with the proof, we remark that the only part of the proof where we distinguish between directed and circular bidirected graphs is the very end where we apply Theorems 12.3.5 and 12.3.6, respectively, and hence get a small additive distinction between the respective bounds.

Let D be either a directed graph or a circular bidirected graph. Let (T, φ) be a cycle-decomposition of D of width $k := \mathrm{cycw}(D)$. So for every edge $e \in E(T)$, the induced cut C_e has cycle-porosity at most k.

Let D_e be the (bi)directed graph obtained by adding a subdivision vertex to each edge of C_e and let S_e be the set of all these subdivision vertices. Since the cycle-porosity is at most k, every union of vertex-disjoint cycles hits at most k vertices of S_e. Applying Theorem 12.2.7 if D_e is a directed graph and Theorem 12.2.3 if D_e is bidirected to the vertex set S_e in D_e provides a set of at most k vertices in $V(D_e)$ that hits all cycles in D_e containing a vertex of S_e. This gives a set Y_e of at most k vertices in $V(D)$ that hits all cycles in D containing an edge of C_e (by possibly replacing vertices of S_e by their neighbours). We now use these sets Y_e to describe a winning strategy for $3k$ cops.

Set $C_0 := Y_e$ as the initial position of the cops for an arbitrary edge $e \in E(T)$. Note that $|C_0| = k \leqslant 3k$ by assumption. Now suppose that after the i-th round, there exists an edge $f \in E(T)$ with $C_i = Y_f$. Then $R_i \subseteq \varphi(L(T_i))$ for a component T_i of $T - f$. If T_i contains only a single vertex t, then the $3k$ cops catch the robber after $i + 1$ rounds by setting $C_{i+1} := C_i \cup \{\varphi(t)\}$ which has size at most $k + 1 \leqslant 3k$.

Otherwise, let f_1 and f_2 be the two edges of T_i incident with f. We set $C_{i+1} := C_i \cup Y_{f_1} \cup Y_{f_2}$. Then the robber must choose a strong component R_{i+1} of $D - C_{i+1}$ that is contained in $\varphi(L(T_{i+1}))$ for a component T_{i+1} of $T_i - f_1 - f_2$ as $C_i \subseteq C_i \cap C_{i+1}$.

Let $j \in \{1,2\}$ with the property that f_j is incident with T_{i+1}. Set $C_{i+2} := Y_{f_j}$ and $T_{i+2} := T_{i+1}$, which ensures that $R_{i+2} \subseteq \varphi(L(T_{i+2}))$ since $C_{i+2} \subseteq C_{i+1} \cap C_{i+2}$. Note that $|C_{i+1}|, |C_{i+2}| \leqslant 3k$ and that T_{i+2} is a proper subtree of T_i.

As D and hence T is finite, the subtrees T_i become smaller as i increases. This implies that the above described strategy indeed enables $3k$ cops to catch the robber on D.

If D is a directed graph, we thus have $\operatorname{dtw}(D) \leqslant 3(3k) - 2 = 9k - 2 = 9\operatorname{cycw}(D) - 2$ by Theorem 12.3.5, as desired. If D is a circular bidirected graph, we have $\operatorname{btw}(D) \leqslant 3(3k) + 2 = 9k + 2 = 9\operatorname{cycw}(D) + 2$ by Theorem 12.3.6, as desired. □

PART IV

Coarse graph theory

CHAPTER 13

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

This chapter is based on [ADEFJKW23] and joint work with Sandra Albrechtsen, Reinhard Diestel, Ann-Kathrin Elm, Eva Fluck, Paul Knappe and Paul Wollan.

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.

13.1 Introduction

13.1.1 Fat minors and quasi-isometries

Following Gromov's ideas on coarse geometry [G93] into the realm of graphs, Georgakopoulos and Papasoglu [GP25] 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. [L17b, 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 13.1.1 ([GP25, 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 $X \in \mathcal{X}$.*

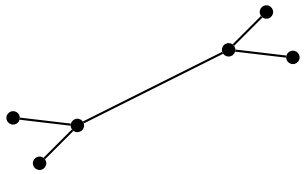

Figure 13.1: The 'wrench'/'spanner' graph W.

Here, 'K-fat minors' are a metric variant of minors: Roughly speaking, a K-fat X minor is an X minor with additional distance constraints: its branch sets and branch paths are pairwise at least K apart, except for incident vertex-edge pairs (see Section 13.1.4 for the formal definition).

Georgakopoulos and Papasoglu [GP25] 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 [AJKW25] for a simpler proof. Moreover, Albrechtsen, Jacobs, Knappe and Wollan [AJKW25] solved the case $\mathcal{X} = \{K^4\}$. In contrast to these positive results, Davies, Hickingbotham, Illingworth and McCarty [DHIM24] showed that Conjecture 13.1.1 is false in general.

In this chapter, we establish Conjecture 13.1.1 for three further small cases but in a stronger form:

Theorem 32. *Conjecture 13.1.1 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 13.1. In all these cases, Conjecture 13.1.1 even holds true when replacing 'fat minors' by 'quasi-geodesic topological minors'.*

We refer the reader to Theorems 33 to 35 for the respective functions f.

In Theorem 32, 'quasi-geodesic topological minors' are a metric version of topological minors, whose model in G is 'quasi-geodesically' embedded (see Section 13.1.4 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 13.3.14).

Our work on Theorem 32 was independent of Conjecture 13.1.1 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 13.1.1, through [GP25]. Our approach is not based on quasi-isometries, but on the notion of graph-decompositions, which we describe in what follows.

13.1.2 Graph-decompositions

Graph-decompositions [DJKK26] 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

[1] Georgakopoulos and Papasoglu prove the more general case $\mathcal{X} = \{K_{1,m}\}$ in a second version of their paper [GP25].

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 [DJKK26] as well as the study of local separations [JKK25] and of locally chordal graphs [AKK25].

In this chapter, we present a width-notion for graph-decompositions such that a graph G has an H-decomposition of small width if and only if H resembles the large-scale structure of G. The naive approach defines the 'width' of a graph-decomposition analogously to tree-width, that is, as the minimal cardinality of a bag of the decomposition (minus 1). However, the respective '$\mathcal{H}$-width', the minimal width of an $\mathcal{H}$-decomposition for a given graph class $\mathcal{H}$, does not yield a meaningful extension of tree-width: if a minor-closed class $\mathcal{H}$ of graphs has bounded tree-width, then every graph of small $\mathcal{H}$-width has small tree-width itself, and if $\mathcal{H}$ has unbounded tree-width, then the $\mathcal{H}$-width of every graph is at most 2 (minus 1) [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 [L17a]).

13.1.3 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

[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].

of trees H by Berger and Seymour [BS24, (4.1)]. Here, we extend their observation to arbitrary graphs H.

Proposition 13.1.2. *There exist functions $g : \mathbb{N}^2 \to \mathbb{N}^2$ and $h_1, h_2 : \mathbb{N}^2 \to \mathbb{N}$ such that the following holds for all graphs G, H:*

(i) *If G admits an honest H-decomposition of radial width r_0 and radial spread r_1, then G is $g(r_0, r_1)$-quasi-isometric to H.*

(ii) *If G is (L, C)-quasi-isometric to H, then G admits an honest H-decomposition of radial width $h_1(L, C)$ and radial spread $h_2(L, C)$.*

We refer the reader to Section 13.3.3 for the detailed statement on the functions g and h_1, h_2.

The correspondence given by Proposition 13.1.2 paves the way for a new proof method towards Conjecture 13.1.1: one through the graph-theoretic construction of a suitable graph-decomposition. We follow this approach in this chapter and further demonstrate its power and versatility in [AJKW25].

Proposition 13.1.2 in particular implies that we reached our initial goal from the beginning of Section 13.1.2 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.

13.1.4 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 13.1.1, 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 13.3.14). Hence, as fat minors form an obstruction to small radial width (Lemma 13.3.15), so do quasi-geodesic topological minors (Proposition 13.3.16).

Our main result, Theorem 32, 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 33 to 35, from which we then deduce Theorem 32.

[3]Note that in metric spaces this property is often called $(c, 0)$-quasi-geodesic. In [BR22] it is called c-multiplicative.

Theorem 33 (Radial path-width)**.** *Let $k \in \mathbb{N}$. If a connected graph G contains no $(\geqslant k)$-subdivision of K^3 as a geodesic subgraph and no $(\geqslant 3k)$-subdivision of $K_{1,3}$ as a 3-quasi-geodesic subgraph, then G admits an honest decomposition modelled on a path P of radial width at most $18k+2$ and radial spread at most $18k+1$.*

Moreover, P is $(1, 18k+2)$-quasi-isometric to G.

Theorem 34 (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 35 (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 32 given Theorems 33 to 35. 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 13.3.14 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 32 from Theorems 33 to 35 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 13.3.8 yields the desired $f(K)$-quasi-isometry from G to H. □

13.1.5 Sketch of the proofs

Let us give a brief overview of the proof techniques for Theorems 33 to 35. For the proof of Theorem 33 (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 33 immediately generalizes to Theorem 34 (radial cycle-width). The proof of Theorem 35 (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

[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 35 in Section 13.6 for details).

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.

13.1.6 An open conjecture about quasi-isometries to trees

For the special case of $\mathcal{X} = \{K^3\}$, Georgakopoulos and Papasoglu answered Conjecture 13.1.1 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 33 to 35 yield another natural guess for such an obstruction – long quasi-geodesic cycles:

Conjecture 13.1.3. *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 13.1.2 and [BS24, (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 13.B we show that the proof technique described in Section 13.1.5 for Theorems 33 and 34 does not work for Conjecture 13.1.3.

In Section 13.A, we show that the strengthening of Conjecture 13.1.1 for quasi-geodesic topological minors analogous to Conjecture 13.1.3 fails earlier than Conjecture 13.1.1 itself.

13.1.7 Structure of the chapter

Section 13.2 collects some basic definitions. In Section 13.3 we introduce the notion of graph-decompositions and their radial width and spread, and show how they are related to quasi-isometries, proving Proposition 13.1.2. We also prove in Section 13.3.4 that quasi-geodesic topological minors yield fat minors and are an obstruction to small radial width. The following three Sections 13.4 to 13.6 contain the proofs of Theorems 33 to 35, respectively. We finally discuss further problems resulting from our work in Sections 13.A and 13.B.

13.2 Preliminaries

We mainly follow the notations from [D25]. 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 $d_G(v,u)$ for the distance of the two vertices v and u in G. For two sets U and U' of vertices of G, we write $d_G(U,U')$ 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 $d_G(v,w)\leqslant k$ 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 $U\subseteq B_G(v,k)$.

13.3 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 13.1.2. Finally, we define quasi-geodesic topological minors and show that they yield fat minors, which indeed form obstruction to small radial width.

13.3.1 Graph-decompositions

Let us recall the notion of graph-decompositions:

Definition 13.3.1 (Graph-decomposition [DJKK26])**.** 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 $hh'\in H$.

[5] In [DJKK26] 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.

If G is connected and every bag of $(H, \mathcal{V})$ is non-empty, then it follows from Definition 13.3.1 that H has to be connected as well. More generally, we have the following lemma:

Lemma 13.3.2. *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.*

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 13.3.3 for tree-decompositions was first proven in [DJKK25, Lemma A.1].

Lemma 13.3.3. *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 $d_G(v, w) \leqslant r$, and we fix a shortest v–w path P in G. Then every vertex $p \in P$ satisfies $d_G(v, p) \leqslant r$ as witnessed by P. In particular, every node $h' \in H$ with $V_{h'} \cap V(P) \neq \emptyset$ satisfies $v \in V'_{h'} = B_G(V_{h'}, r)$, and hence H_P is a subgraph of H'_v. But now $h \in V(H_P)$ as P meets the vertex $w \in V_h$, and H_v is a subgraph of H_P as $v \in V(P)$. So since H_P is connected by Lemma 13.3.2, there exists a h–$V(H_v)$ path in H_P and hence in H'_v, as desired. □

13.3.2 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 13.3.4 (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

$$\operatorname{radw}(H, \mathcal{V}) := \max_{h \in V(H)} \operatorname{rad}(G[V_h]).$$

Note that if $G[V_h]$ is disconnected for some $h \in H$, then $\operatorname{radw}(H, \mathcal{V}) = \infty$. The *outer-radial width* of $(H, \mathcal{V})$ is $\max_{h \in V(H)} \operatorname{rad}_G(V_h)$. 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

$$\operatorname{radw}_{\mathcal{H}}(G) := \min \{\operatorname{radw}(H, \mathcal{V}) \mid (H, \mathcal{V}) \text{ is a graph-decomposition of } G \text{ with } H \in \mathcal{H}\}.$$

Note that the (inner-)radial $\mathcal{H}$-width will always be at most $\operatorname{rad}(G)$ for a connected graph G if $\mathcal{H}$ contains at least one non-empty graph. The outer-radial $\mathcal{H}$-width is defined analogously.

For classes $\mathcal{H}$ of graphs that we frequently use in this chapter, we name the radial $\mathcal{H}$-width as in Table 13.1.

Table 13.1: 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 13.A).

$\mathcal{H}$ **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; L17a] with applications to tree-spanners and routing problems.

Definition 13.3.5 (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 13.3.6, 13.3.9 and 13.3.10. Hence, we will usually omit the ‘inner’ when talking about (inner)-radial width and spread.

Berger and Seymour [BS24, (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 13.3.6. *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 13.3.3, $(H, \mathcal{V}')$ 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 $d_{G[V'_h]}(x', x) = d_G(x', x) \leqslant k$. Thus, we have $d_{G[V'_h]}(z_h, x) \leqslant d_{G[V'_h]}(z_h, x') + d_{G[V'_h]}(x', x) \leqslant k + k = 2k$. This implies that each $G[V'_h]$ has radius at most $2k$, and hence $(H, \mathcal{V}')$ has (inner-)radial width at most $2k$.

To see that $(H, \mathcal{V}')$ has (inner-)radial spread at most $2kr$, consider any vertex v of G and any node h of H such that $v \in V'_h$. By the definition of V'_h, there exists a vertex $u \in V_h$ such that $d_G(v, u) \leqslant k$. Let $P = p_0 \dots p_n$ be a shortest v–u path in G, in particular, $||P|| = n \leqslant k$. By Lemma 13.3.2, 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 $2rk$. □

13.3.3 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 13.1.2. For this, let us recall the definition of quasi-isometries in the case of graphs.

Definition 13.3.7 (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. [GP25, 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 13.3.8. *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 13.1.2 into the following two lemmas.

Lemma 13.3.9. *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 $h_0 \coloneqq h$ and $h_{\ell+1} \coloneqq h'$. For every $i \in \{0, \dots, \ell\}$, let P_i be a shortest h_i–h_{i+1} path in H. As $h_i, h_{i+1} \in V(H_{v_i})$ for all $i \in \{0, \dots, \ell\}$ and $(H, \mathcal{V})$ has outer-radial spread r_1, all P_i have length at most $2r_1$. Thus, $h_0P_0h_1 \dots h_\ell P_\ell h_{\ell+1}$ is an h–h' walk in H of length at most $2r_1(\ell + 1)$, as desired. □

Lemma 13.3.10. *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, $d_G(\varphi(h_0), \varphi(h_1))$ 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, $d_G(\varphi(h'_0), \varphi(h'_1))$ 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

$$d_G(v, \varphi(h)) \leqslant d_G(v, \varphi(h')) + d_G(\varphi(h), \varphi(h')) \leqslant r' + M(mr' + \lceil (m+a)/2 \rceil) + A.$$

Thus, every $v \in V_h$ has distance at most $r' + M(mr' + \lceil (m+a)/2 \rceil) + A$ from $\varphi(h)$. This yields the desired outer-radial width.

To obtain the moreover-part, let us investigate the above equation in more detail. Any shortest v–$\varphi(h')$ in G is contained in $B'_h \subseteq V_h$. Fix a shortest h'–h path Q and replace each edge h_0h_1 of Q by a shortest $\varphi(h_0)$–$\varphi(h_1)$ path $Q_{h_0h_1}$ in G to obtain a $\varphi(h)$–$\varphi(h')$ walk Q' in G. If $r' \geqslant M + A$, the path $Q_{h_0h_1}$ is contained in $B'_{h_0} \subseteq V_h$. Thus, the (inner-)radial width is already at most $r' + M(mr' + \lceil (m+a)/2 \rceil) + A$. □

Proof of Proposition 13.1.2. Use Lemma 13.3.9 for (i) and Lemma 13.3.10 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, (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 13.3.11. *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 13.3.11 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.

13.3.4 Quasi-geodesic topological minors

In this section, we study quasi-geodesic topological minors as obstructions to small radial width. We show in Lemma 13.3.14 that quasi-geodesic topological minors are also an instance of the more general 'fat minors' used in Conjecture 13.1.1. Fat minors are obstructions to small radial width (Lemma 13.3.15), which then implies that quasi-geodesic topological minors are also such obstructions (Proposition 13.3.16).

Let us recall the definition of $(\geqslant k)$-subdivisions T_kX and c-quasi-geodesity from the introduction.

Definition 13.3.12 (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 13.3.13 (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 $d_G(u,v) \leqslant c \cdot d_X(u,v)$. We call X *geodesic* if it is 1-quasi-geodesic.

Let us now turn to the question how quasi-geodesic topological minors relate to fat minors, the obstruction investigated by Georgakopoulos and Papasoglu in Conjecture 13.1.1. 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 $d_G(Y,Z) \geqslant K$ 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 13.3.14. *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 least $(k+1) - 2\lceil (k+1)/4 \rceil \geqslant k/2 - 1$. We then let $\mathcal{V}$ be the set of all V_x and $\mathcal{E}$ be the set of all E_e. Since X_G is a T_kX, this construction directly yields $d_{X_G}(Y,Z) \geqslant k/2 - 1$

[6] Note that we here deviate from the definition of model from [D25]. 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 [D25].

for every two distinct $Y, Z \in \mathcal{V} \cup \mathcal{E}$ unless $Y = E_e$ and $Z = V_x$ for some vertex $x \in V(X)$ incident to $e \in E(X)$, or vice versa. Since X_G is a c-quasi-geodesic subgraph of X, we then obtain $d_G(Y, Z) \geqslant d_{X_G}(Y, Z)/c \geqslant \frac{k-2}{2c}$; so X_G is indeed a K-fat model of X in G. □

Lemma 13.3.15. *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 13.3.2.

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 13.3.2, 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 13.3.16. *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.* □

13.4 Radial path-width

In this section we prove Theorem 33, which we restate here for convenience.

Theorem 33. *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 33.

Theorem 13.4.1. *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 13.4.1 indeed implies Theorem 33. 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 13.2).

Lemma 13.4.2. *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 13.3.10, 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 13.3.10 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 $k := \lfloor (j - i)/2 \rfloor$. □

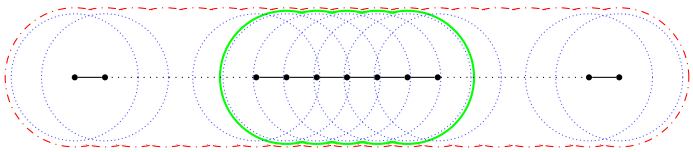

Figure 13.2: A decomposition modelled on the black path using the blue balls as given by Lemma 13.4.2 for $r = 2$, $c = 1$. The bag corresponding to the centre vertex is green.

Proof of Theorem 33 given Theorem 13.4.1. 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 13.4.1, 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 13.4.2 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 13.4.1. 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 13.4.3, 13.4.4 and 13.4.6 below. They are stated in a slightly more general form than needed for the proof of Theorem 13.4.1, which enables us to use them later also in the proofs of Theorems 34 and 35.

The first lemma, Lemma 13.4.3 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 13.4.4, 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 13.3). The third lemma, Lemma 13.4.6, 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 13.4.3. *Let G be a graph, and let X be a c-quasi-geodesic subgraph of G for some $c \in \mathbb{N}$. If P is a shortest v–X path in G for some vertex $v \in G$, then $X \cup P$ is $(2c+1)$-quasi-geodesic in G.*

Proof. We have to show that, for every two vertices u and w of $X \cup P$, the distance of u and w in $X \cup P$ is at most $(2c+1)$ times their distance in G. Since X is c-quasi-geodesic in G and P is geodesic in G, it is (by symmetry) enough to consider the case where u is a vertex of P and w is a vertex of X. Let x be the endvertex of P in X. Since P is a shortest v–X path in G, uPx is a shortest u–X path in G and hence $d_P(u,x) = d_G(u,x) \leqslant d_G(u,w)$ as $w \in X$. We then have $d_X(x,w) \leqslant c \cdot d_G(x,w) \leqslant c \cdot (d_G(x,u) + d_G(u,w))$, where the first inequality follows from X being c-quasi-geodesic in G while the second one applies the triangle inequality. Again using the triangle inequality, we can then combine these inequalities to

$$\begin{aligned} d_{X\cup P}(u,w) &\leqslant d_{X\cup P}(u,x) + d_{X\cup P}(x,w) = d_P(u,x) + d_X(x,w) \\ &\leqslant d_G(u,x) + c \cdot (d_G(x,u) + d_G(u,w)) \leqslant (2c+1) \cdot d_G(u,w), \end{aligned}$$

which shows the claim. □

Lemma 13.4.4. *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 13.3 for a sketch of the situation in Lemma 13.4.4.

The proof of Lemma 13.4.4 builds on the study of how cycles interact with a given separation of the graph. More formally, we have the following preparatory lemma.

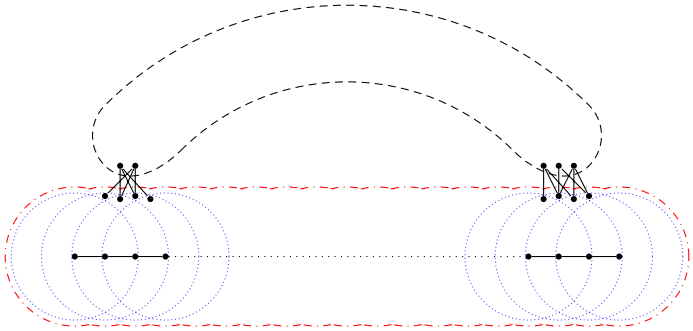

Figure 13.3: The setting of Lemma 13.4.4.

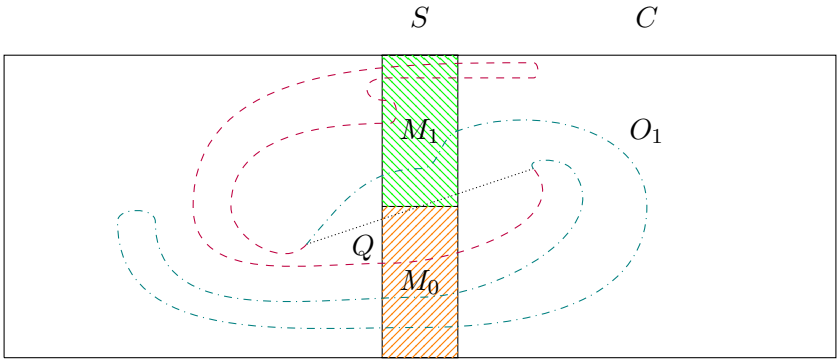

Figure 13.4: The setting of Lemma 13.4.5. Here, O_2 is the cycle containing Q (black, dotted line) and the purple, dashed part of O_1, and O_3 is the cycle containing Q and the teal, dash dotted part of O_1.

Lemma 13.4.5. *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 13.4 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 13.4.4. Note that, as $n > m_0 + m_1$, we in particular have $n - m_1 > m_0$. We set

$$M_0 := N_G(C) \cap \bigcup_{i=0}^{m_0} B_i \text{ and } M_1 := N_G(C) \cap \bigcup_{i=n-m_1}^{n} B_i.$$

As C has no neighbours in $\bigcup_{i=m_0+1}^{n-m_1-1} B_i$, the neighbourhood of C is equal to $M_0 \cup M_1$ and C is a component of $G - M_0 - M_1$.

First, we note that M_0 and M_1 are disjoint. Indeed, if $u \in B_{j_0}$ for some $j_0 \leqslant m_0$ and $v \in B_{j_1}$ for some $j_1 \geqslant n - m_1$, then using the triangle inequality we obtain

$$n-m_0-m_1 \leqslant d_P(p_{j_0}, p_{j_1}) = d_G(p_{j_0}, p_{j_1}) \leqslant d_G(p_{j_0}, u)+d_G(u, v)+d_G(v, p_{j_1}) = d_G(u, v)+2r. \tag{13.1}$$

So $d_G(u, v) \geqslant n - m_0 - m_1 - 2r = m > 0$ by assumption. In particular, u and v are distinct vertices, and thus M_0 and M_1 are disjoint.

We now define, for a cycle O in G, the number $W(O)$ as the number of M_0–M_1 paths in O which have an inner vertex in C.

The remainder of the proof now consists of three steps. First, we show that a cycle O with $W(O) \neq 0$ has length at least $2m$. Second, we find a cycle O with odd $W(O)$, and lastly we show that a shortest such cycle is geodesic in G.

Claim 1. *Every cycle O in G with $W(O) \neq 0$ has length at least $2m$.*

Proof of the claim. 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 (13.1), 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 of the claim. 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 of the claim. 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 13.4.5 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 13.4.6. *Let U be a set of vertices of a connected graph G, and let $p_0 \in V(G) \setminus U$ have maximal distance from U in G. Let $P = p_0 \dots p_n$ be a shortest p_0–U path in G. Given $r \in \mathbb{N}$, we write $B_i := B_G(p_i, r)$ for every $p_i \in P$. Suppose that some component C of $G - \bigcup_{i=0}^{n} B_i$ is disjoint from U and satisfies $N_G(C) \subseteq \bigcup_{i=0}^{\ell} B_i$ for some $\ell \leqslant n$. Then $d_G(v, P) \leqslant 2r + \ell$ for every $v \in C$.*

Proof. Let v be an arbitrary vertex of C, and let Q be a shortest v–U path in G with endvertex $u \in U$. As $V(C) \cap U = \emptyset$, the path Q intersects $\bigcup_{i=0}^{n} B_i$. Let q be the first vertex on Q in $\bigcup_{i=0}^{n} B_i$, and let j be the smallest index of a ball B_j with $q \in B_j$. Note that $j \leqslant \ell$ since $N_G(C) \subseteq \bigcup_{i=0}^{\ell} B_i$.

Since P is a longest geodesic path in G which ends in U, we have that $n = ||P|| \geqslant ||Q|| = d_G(v, u) = d_G(v, q) + d_G(q, u)$, where the last equality follows from Q being geodesic in G. Additionally, since $p_n \in U$, we have $d_G(p_j, p_n) \leqslant d_G(p_j, u) \leqslant d_G(p_j, q) + d_G(q, u) = r + d_G(q, u)$. These two inequalities combine to

$$d_G(v, q) \leqslant n - d_G(q, u) \leqslant n - (d_G(p_j, p_n) - r).$$

All in all, we get

$$d_G(v, P) \leqslant d_G(v, p_j) \leqslant d_G(v, q) + r \leqslant n - (d_G(p_j, p_n) - r) + r = n - (n - j) + 2r \leqslant 2r + \ell,$$

where the last inequality follows from $j \leqslant \ell$. This concludes the proof of the claim. □

With all the previous lemmas at hand, we are now ready to prove Theorem 13.4.1.

Proof of Theorem 13.4.1. 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 of the claim. 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 13.4.3, $P \cup Q$ 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 of the claim. Suppose for a contradiction that there is such a component C. Applying Lemma 13.4.4 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 of the claim. If $N_G(C) \subseteq \bigcup_{i=0}^{3k} B_i$, then the claim follows by applying Lemma 13.4.6 with $U = \{p_n\}$. Otherwise, $N_G(C) \subseteq \bigcup_{i=n-3k}^{n} B_i$, and the claim follows by applying Lemma 13.4.6 with $U = \{p_0\}$. ■

By Claims 1 to 3, every vertex of G has distance at most $9k$ from P, as desired. □

13.5 Radial cycle-width

In this section we build on our result on radial path-width, Theorem 33, and use a small adaptation of its proof to address radial cycle-width by proving Theorem 34. More precisely, Theorem 34 follows from the lemmas that we have already shown in Section 13.4. Let us restate Theorem 34 here for convenience.

Theorem 34. *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 34.

Theorem 13.5.1. *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 13.5.1 implies Theorem 34.

Proof of Theorem 34 given Theorem 13.5.1. 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 13.5.1 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 13.4.2 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 13.5.1.

Proof of Theorem 13.5.1. Let G be a connected graph that contains no $T_{3k}K_{1,3}$ as a 3-quasi-geodesic subgraph. Applying Theorem 13.4.1 to G yields that G either contains a geodesic path P such that $C := P$ is as desired, or G contains a $T_{4k+1}K^3$ as a geodesic subgraph. Since we are done in the first case, we may thus assume that G contains a $T_{4k+1}K^3$ as a geodesic subgraph, which means that there exists a geodesic cycle C in G of length at least $3 \cdot (4k+2) = 12k+6$. Set $G_C := G[B_G(C, 3k)]$. If $G = G_C$, then we are done.

Otherwise, there exists a neighbour v of $V(G_C)$ in G since G is connected. Let P be a shortest v–C path in G, and let u be its endvertex in C. Since $v \in N_G(G_C)$, it has distance $3k+1$ from C, and thus P has length $3k+1$. Let $Q := B_C(u, 3k+1)$ be the subpath of C of length $6k+2$ which contains u as its 'middle' vertex. Note that Q is actually a path as C has length at least $12k+6 > 6k+2$. Since P is a path from v to C which ends in u, the paths P and Q only meet in u. Hence, $P \cup Q$ is a $T_{3k}K_{1,3}$ in G.

To conclude the proof and obtain the desired contradiction to our assumptions on G, it remains to show that $P \cup Q$ is 3-quasi-geodesic in G. Since Q is a subpath of C with $||Q|| = 6k+2 \leqslant \frac{1}{2}(12k+6) \leqslant \frac{1}{2}||C||$, we have $d_Q(u,v) = d_C(u,v)$ for every two vertices $u, v \in Q$. Thus, Q is geodesic in G, since C is geodesic in G. By Lemma 13.4.3, $P \cup Q$ then is a 3-quasi-geodesic subgraph of G, as desired. □

13.6 Radial star-width

In this section we prove Theorem 35, which we restate here for convenience.

Theorem 35. *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 13.1. 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 35 for details).

Before we start with the proof of Theorem 35, we first show an auxiliary lemma. Recall that in the proof of Theorem 33 we used Lemma 13.4.4 as a tool to show that a graph contains a geodesic T_kK^3. Similarly, we will use the following lemma in the proof of Theorem 35 to show that a graph contains a 3-quasi-geodesic $T_{3k}W$.

Lemma 13.6.1. *Let $P = p_0 \dots p_n$ be a geodesic path in a graph G. Given two vertices $u, v \in G$, let Q be a shortest u–P path in G with endvertex $q \in P$, and let R be a shortest v–$(P \cup Q)$ path in G with endvertex $r \in P \cup Q$. Suppose that at least one of the following two conditions hold:*

(i) $r \in P$ and $d_G(q, r) \geqslant 4 \max\{d_G(u, P), d_G(v, P)\}$;

(ii) $r \in Q$ and $d_G(q, r) \geqslant 4 \max\{d_G(v, Q), d_G(p_0, q), d_G(p_n, q)\}$.

Then $P \cup Q \cup R$ is a 3-quasi-geodesic subgraph of G.

Proof. We first consider the case that $r \in P$. By Lemma 13.4.3, we have that $P \cup Q$ and $P \cup R$ are each 3-quasi-geodesic subgraphs of G. Hence, in order to show that $X := P \cup Q \cup R$ is 3-quasi-geodesic, it suffices to consider vertices $x, y \in X$ with $x \in Q$ and $y \in R$.

Since P is geodesic, we have that $d_X(q, r) = d_G(q, r) \leqslant d_G(q, x) + d_G(x, y) + d_G(y, r)$ and thus

$$\begin{aligned} d_G(x, y) &\geqslant d_G(q, r) - d_G(x, q) - d_G(y, r) \\ &\geqslant 4 \max\{d_G(u, P), d_G(v, P)\} - d_G(u, P) - d_G(v, P) \\ &\geqslant 2 \max\{d_G(u, P), d_G(v, P)\} \geqslant 2 \max\{d_G(x, q), d_G(y, r)\}. \end{aligned}$$

Therefore, making again use of $d_G(q, r) \leqslant d_G(q, x) + d_G(x, y) + d_G(y, r)$, we obtain

$$\begin{aligned} d_X(x, y) &= d_X(x, q) + d_X(q, r) + d_X(r, y) = d_G(x, q) + d_G(q, r) + d_G(r, y) \\ &\leqslant 2d_G(x, q) + 2d_G(y, r) + d_G(x, y) \leqslant 3d_G(x, y). \end{aligned}$$

For the second case, assume that $r \in Q$. Similar as before, we find that $P \cup Q$ and $Q \cup R$ are each 3-quasi-geodesic subgraphs of G. Hence, we only need to consider vertices $x, y \in X$ with $x \in P$ and $y \in R$. As before, we find that

$$d_G(x, y) \geqslant d_G(q, r) - d_G(x, q) - d_G(y, r) \geqslant 2 \max\{d_G(v, Q), d_G(p_0, q), d_G(p_n, q)\},$$

and thus

$$d_X(x, y) = d_G(x, q) + d_G(q, r) + d_G(r, y) \leqslant 2d_G(x, q) + 2d_G(y, r) + d_G(x, y) \leqslant 3d_G(x, y). \quad \square$$

Now we prove Theorem 35. We remark that we did not optimise the bounds on the radial width and radial spread.

Proof of Theorem 35. 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 13.3.8 and 13.3.9, there exists an (L, C)-quasi-isometry from G to S such that L and C depend on k only. Applying [NSS25b, p. 1.3] (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 13.3.8, 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 of the claim. 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 13.4.4 to obtain a geodesic cycle in G of length at least $2 \cdot (2k + 2) > 3k + 3$ which contradicts our assumption that G does not contain a geodesic $T_k K^3$.

Thus, for every component C of $G - G_P$, either $N_G(C) \subseteq \bigcup_{i=0}^{23k+4} B_i$ or $N_G(C) \subseteq \bigcup_{i=n-23k-4}^{n} B_i$ (this includes in particular the case that $G = G_P$). Then by Lemma 13.4.6 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 13.4.2 applied to the geodesic path P and to $r = 29k + 4$ we obtain an honest P-decomposition of G of radial width at most $2 \cdot (29k + 4) + 1 = 58k + 9$ and radial spread at most $2 \cdot (29k + 4) + 1 = 58k + 9$. ■

By Claim 1 (and since P is a path and hence a subdivided star) we may assume that there is some s with $23k + 5 \leqslant s \leqslant n - 23k - 5$ such that some vertex of $G - G_P$ has a neighbour in B_s. We fix s for the rest of the proof. The following two claims now show that every component of $G - G_P$ can only attach to G_P either close to the start of P or close to the end of P or close to p_s.

Claim 2. *If a vertex of $G - G_P$ has a neighbour in a ball B_t, then $t \leqslant 3k$ or $s - 12k - 3 \leqslant t \leqslant s + 12k + 3$ or $n - 3k \leqslant t$.*

[7] Note that we use here that subdivided stars admit path-decompositions into bags of size 3 and thus have path-width 2.

Proof of the claim. 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 13.6.1 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 $G - G_P$, we have that either $N_G(C) \subseteq \bigcup_{i=0}^{3k} B_i$ or $N_G(C) \subseteq \bigcup_{i=s-12k-3}^{s+12k+3} B_i =: D_s$ or $N_G(C) \subseteq \bigcup_{i=n-3k}^{n} B_i$.*

Proof of the claim. 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 13.4.4 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 13.4.4 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 13.4.6 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 of the claim. 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 13.4.6 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 13.4.2, 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 $\mathcal{C}$ be the set of components that contain a vertex whose distance from P is at least $24k+5$, and let $\mathcal{C}'$ be the set of all other components of $G - V''$. We now add all vertices in $V'' \setminus V'$ and all components from $\mathcal{C}'$ to the bag V_s' that is indexed by p_s. More precisely, we define

$$V_s := V_s' \cup (V'' \setminus V') \cup \{V(C) \colon C \in \mathcal{C}'\}.$$

Further, we set $V_i := V_i'$ for all $i \neq s \in [n]$. Let also $V''' := V'' \cup V_s$ and $G''' := G[V''']$.

Claim 5. *$(P, \mathcal{V})$ is a decomposition of G''' of radial width at most $36k+7$ and radial spread at most $48k+8$ and every vertex in V_s has distance at most $36k+7$ from p_s in $G[V_s]$.*

Proof of the claim. Since $N_G(v) \cap V'$ is contained in $V_s' \subseteq V_s$ for every vertex $v \in N_G(G')$ by Claim 4, and because $(P, \mathcal{V}')$ is a decomposition of G', it follows that $(P, \mathcal{V})$ is a decomposition of G'''. Its radial spread is at most $48k+8$ by Claim 4 and because each vertex in $G''' - V'$ is only contained in V_s. Moreover, every part $G[V_i]$ with $i \neq s$ has radius at most $24k+5$ by Claim 4, so it remains to consider $G[V_s]$.

Let $v \in V_s$ be given. If $v \in V_s'$, then $d_{G[V_s]}(v, p_s) \leqslant 24k+5$ by Claim 4. Otherwise, v has distance at most $24k+4$ from some vertex p_i of P, where $s - 12k - 3 \leqslant i \leqslant s + 12k + 3$ by Claim 2, and hence

$$d_{G[V_s]}(v, p_s) \leqslant d_{G[V_s]}(v, p_i) + d_{G[V_s]}(p_i, p_s) \leqslant (24k+4) + (12k+3) = 36k+7$$

where we used that $p_i P p_s \subseteq G[V_s'] \subseteq G[V_s]$ by the definition of V_s'. ■

We now show that all components of $G - V'''$, i.e. those in $\mathcal{C}$, are path-like.

Claim 6. *Let C be a component in $\mathcal{C}$. Then there is an honest decomposition of $G[V(C) \cup V_s]$ modelled on a path Q of radial width at most $54k+8$ and radial spread at most $12k+2$ such that the bag corresponding to the last node of Q contains V_s, and V_s only meets bags assigned to nodes of distance at most $6k$ to the last node of Q.*

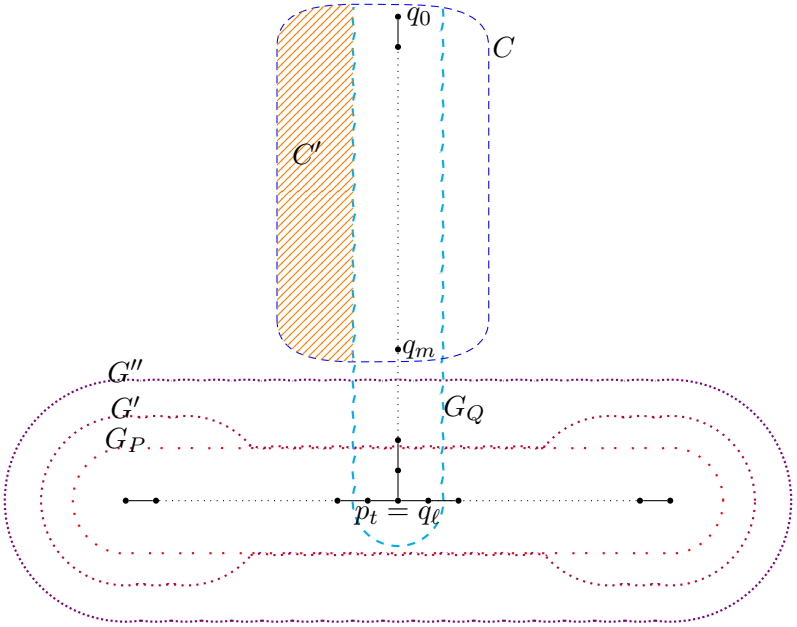

Figure 13.5: The setting of the proof of Claim 6.

Proof of the claim. The reader may look at Figure 13.5 to follow the proof more easily. Let q_0 be a vertex of C of maximal distance from P. Since $C \in \mathcal{C}$ the vertex q_0 has distance at least $24k+5$ from P. Let $Q = q_0 \dots q_\ell$ be a shortest path in G from q_0 to P (so Q is geodesic in G). Note that C is contained in a component of $G - G'$, and thus Q ends in a vertex $p_t = q_\ell$ with $s - 12k - 3 \leqslant t \leqslant s + 12k + 3$ by Claim 2. In particular $3k < t < n - 3k$. Let m be the last index such that $q_m \notin G''$. Then Qq_m is a shortest q_0–$N_G(G'')$-path, and $m = \ell - 18k - 4$.

For every $0 \leqslant i \leqslant \ell$ we define B_i^Q to be the ball in G of radius $3k$ around q_i, and let $G_Q = G[\bigcup_{i=0}^{\ell} B_i^Q]$. We claim that every vertex of $C - G_Q$ has distance at most $12k$ from q_0. To see this, let C' be any component of $C - G_Q$, and let C'' be the component of $G - G_Q$ containing C'. We first show that C'' has no neighbour in $\bigcup_{i=3k+1}^{\ell} B_i^Q$ and that this in particular implies that $C' = C''$.

Towards a contradiction, we first assume that C'' has a neighbour in some B_j^Q with $3k < j < \ell - 12k - 3$. Then $q_j Q$ has length at least $12k + 4 = 4 \cdot (3k+1)$. So Lemma 13.6.1 (ii) applied to $p_{t-3k-1} P p_{t+3k+1}$, $q_{j-3k-1} Q$ 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 13.4.4 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 13.4.6 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 13.4.2, the U_i' are the bags of a Q-decomposition of G_Q of radial spread at most $4 \cdot (3k) + 2 = 12k + 2$; and every vertex in any U_i' has distance at most $6k + 1$ from q_i in $G[U_i']$.

In a second step, for $Q^* = Qq_m$, we define a Q^*-decomposition of $G_Q^* \cup C$ where $G_Q^* := G[\bigcup_{i=0}^{m+3k+1} B_i^Q]$. For that, we let $U_i'' := U_i'$ for $1 \leqslant i \leqslant m$. The neighbourhood of every component of $C - G_Q$ is contained in U_0', so we can just add its vertices to U_0': let U_0'' be the union of U_0' and all vertices of components of $C - G_Q$. By construction U_0'' has radius at most $12k$. Then the U_i'' are the bags of a Q^*-decomposition of $G_Q^* \cup C$. Indeed, every vertex of $C - G_Q$ is contained in U_0'' by definition, and every vertex of $C \cap G_Q$ is contained in G_Q^* since all vertices in $\bigcup_{i=m+3k+2}^{\ell} B_i^Q$ have distance at most $18k + 3$ from P and are thus not contained in C. Moreover, $(Q^*, \mathcal{U}'')$ has radial width at most $12k$ and radial spread at most $12k + 2$.

We now restrict the bags U_i'' to C and add V_s to all U_i'' with $i \geqslant m - 6k$, i.e. we set $U_i := (U_i'' \cap V(C))$ for $i < m - 6k$ and $U_i := (U_i'' \cap V(C)) \cup V_s$ for all $i \geqslant m - 6k$. To verify that $(Q^*, \mathcal{U})$ is a decomposition of $G[V(C) \cup V_s]$, it suffices to show that $N_G(G - C) \subseteq U_m$.

For this, let $v \in N_G(G - C) \subseteq V(C)$ be given, and let w be a neighbour of v in $G - C$. As C is a component of $G - G''$, this means that $w \in G''$. As we have already seen above, no component of $C - G_Q$ has neighbours in G'', so $v \in G_Q$, which in particular implies that $v \in G_Q^*$. If $v \in \bigcup_{i=m-3k-1}^{m+3k+1} B_i^Q$, then $v \in U_m$ by construction. So we may assume that $v \in B_j^Q$ for some $j < m - 3k - 1$. But this implies that q_0 has distance at most

$$d_G(q_0, q_j) + d_G(q_j, v) + d_G(v, w) + d_G(w, P) \leqslant (m - 3k - 2) + 3k + 1 + (18k + 4) = m + 18k + 3 < \ell$$

from P, which contradicts that Q is a shortest q_0–P path.

Hence, $(Q^*, \mathcal{U})$ is a path-decomposition of $G[V(C) \cup V_s]$, and it has radial spread at most $12k + 2$ by construction. Moreover, it has radial width at most $54k + 8$: Let $i \leqslant m$ and let $v \in U_i$. If $i \leqslant m - 6k$, then, since $Q^* \subseteq C$ is a shortest q_0–$N_G(G - C)$ path in G

and $G[U_i']$ has radius at most $12k$, we have $U_i = U_i'$ and hence $G[U_i]$ has radius at most $12k$. Now suppose that $i > m - 6k$. If $v \in V_s$, then v has distance at most $36k + 7$ from p_s by Claim 5. Otherwise, if $v \in U_i \setminus V_s$ then

$$d_{G[U_i]}(v, p_s) < d_{G[U_i]}(v, q_i) + d_{G[U_i]}(q_i, q_m) + d_{G[U_i]}(q_m, p_s) \leqslant 12k + 6k + (36k + 8) = 54k + 8,$$

where we used that $q_i Q q_m \subseteq G[U_i]$ by definition, that $q_m \in N_G(V_s') \subseteq V_s$ by Claim 4, and that every vertex in V_s has distance at most $36k + 7$ from p_s by Claim 5. ■

We now combine these path-decompositions of the components of $G - V'''$ with the path-decomposition $(P, \mathcal{V})$ of G''' to a star-decomposition of G. For this, recall that by Claim 5, $(P, \mathcal{V})$ is a path-decomposition of G''' of radial width at most $36k + 7$ and radial spread at most $48k + 8$. Moreover, the components in $\mathcal{C}$ are precisely the components of $G - V'''$. For every component $C \in \mathcal{C}$, Claim 6 guarantees the existence of a decomposition $(Q_C, \mathcal{U}_C)$ of C modelled on a path Q_C of radial width at most $54k + 8$ and radial spread at most $12k + 2$ such that the bag in $\mathcal{U}_C$ corresponding to the last vertex of Q_C contains V_s.

We obtain a subdivided star S from the disjoint union of P and the Q_C by adding edges from the last vertices of the Q_C to p_s. Now every vertex h of S already has a bag, which we denote by V_h, in exactly one of the path-decompositions $(P, \mathcal{V})$ or $(Q_C, \mathcal{U}_C)$. It is straightforward to check that $(S, \mathcal{V})$ is a star-decomposition of G and that is radial width is at most $54k + 8$. Moreover, its radial spread is at most $(48k + 8) + 1 + 6k = 52k + 9$, as only vertices in V_s may lie in more than one decomposition of the form $(P, \mathcal{V})$ or $(Q_C, \mathcal{V})$. This completes the proof. □

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 35 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.

Acknowledgements

We are thankful to an anonymous reviewer for the detailed feedback that helped us improve the presentation of this work.

13.A Where quasi-geodesic topological minors are not the right obstruction

Davies, Hickingbotham, Illingworth and McCarty showed that Conjecture 13.1.1 is false in general. Hence, the strengthening of Conjecture 13.1.1 to quasi-geodesic topological minors

also has a negative answer. However, this strengthening fails earlier than Conjecture 13.1.1, and in this appendix 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 13.A.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 13.A.1:

Corollary 13.A.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 13.A.1) and by a finite list $\{X_1, \dots, X_\ell\}$ of forbidden topological minors (Corollary 13.A.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 13.1.1 for quasi-geodesic topological minors:

Question 13.A.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 13.A.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 13.A.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 [W37]. The following proposition provides a sufficient criterion for this, via the detour of graph-decompositions as established in Proposition 13.1.2.

Proposition 13.A.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 13.A.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 13.3.16, 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. □

13.B Counterexamples to stronger versions in the star and tree case

A possible strengthening of Theorem 35 and Conjecture 13.1.3 is suggested by the proofs of the path- and cycle-case (see Theorems 13.4.1 and 13.5.1). 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 35 and Conjecture 13.1.3. The star-decomposition constructed in the proof of Theorem 35, 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 13.B.3 for more details).

Example 13.B.1. Let $d, k \geqslant 1$ be integers. Let G be a subdivision of the graph W depicted in Figure 13.1 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 35 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 13.B.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 13.B.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 13.1.3 will have to build on new methods for constructing graph-decompositions of low radial width.

13.B.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 13.B.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 $1 \leqslant c < 2$.

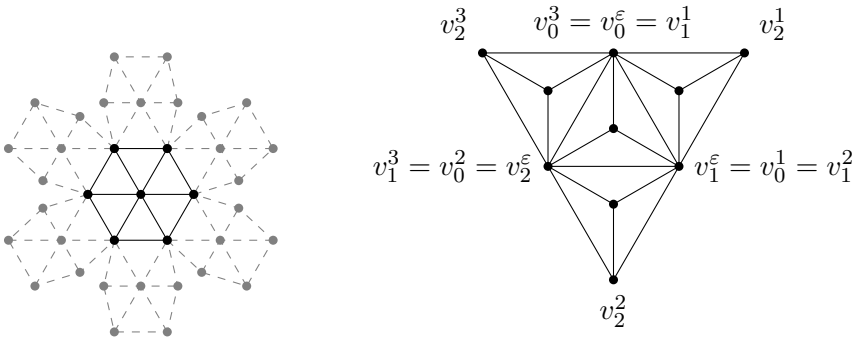

Figure 13.6: Two trees of wheels. The left graph is a tree of wheels of depth 1 for $n = 6$, and the edges that do not belong to the central wheel are dashed. The right graph is a tree of wheels of depth 1 for $n = 3$, and the vertices that are not hubs are labelled.

(ii) ([BS24]) Assume $k \geqslant 2$. The $k \times k$ grid has radial tree-width at least $(k-2)/2$ by Proposition 13.3.16, 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 13.B.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 13.B.3. Let $k, d \geqslant 0$ be integers. Let G be a subdivision of the graph W depicted in Figure 13.1 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 13.B.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.

Example 13.B.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 13.6. 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 13.B.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 $I \subseteq \{1, \dots, n\}$

[8] The proof can be adapted to $n = 3$, but then we get a different bound on c.

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 13.B.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. 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 $v_1^s = v_1^{s2} = \dots = v_1^{st}$ 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 13.B.5 yields $c \geqslant 2(d-(p-1)-(c+1)+1)$ and thus $c \geqslant (2/3) \cdot (d - \lceil q/2 \rceil + 1)$.

[9]The proof can be adapted to $n \geqslant 3$, but then we get a different bound on c.

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 13.B.5. □

CHAPTER 14

A characterisation of the graphs quasi-isometric to K^4-minor-free graphs

This chapter is based on [AJKW25] and joint work with Sandra Albrechtsen, Paul Knappe and Paul Wollan.

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.

14.1 Introduction

14.1.1 Quasi-isometry and fat minors

Gromov's [G93] 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 14.2.3 for the definition).

Following Gromov's idea into the realm of graphs, Georgakopoulos and Papasoglu [GP25] 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 [GP25, 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 14.2.2 for the definition).

Conjecture 36. *[GP25, 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 36 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 36 holds for any (not even necessarily finite) graph X.

Recently, Conjecture 36 has been disproved by Davies, Hickingbotham, Illingworth and McCarty for general graphs X [DHIM24, Theorem 1]. However, Conjecture 36 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 [M05] (see [GP25] 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 [GP25]. 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.

14.1.2 Our results

The main contribution of this chapter is the resolution of Conjecture 36 for $X = K_4$:

Theorem 37. *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 37, 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 36 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 makes it possible to look at Conjecture 36 from a different angle: Graph-decompositions [DJKK26] 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 [DJKK26] as well as the study of local separations [JKK25] and of locally chordal graphs [AKK25]. In [ADEFJKW23, § 3.3] 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 14.2.4 and 14.2.5 for details). To prove Theorem 37, 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 37.

Our techniques additionally yield a new proof of Conjecture 36 for $X = K_4^-$, which is significantly shorter than the original one by Fujiwara and Papasoglu [FP23]:

Theorem 38. *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 38. 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 37 and 38 for arbitrary (finite or infinite) graphs[2]. Thus, Theorems 37 and 38 transfer verbatim to length spaces, a generalisation of geodesic metric spaces [GP25, Observation 2.1].

14.1.3 Applications

Bonamy, Davies, Esperet and Wesolek brought the following consequences of Theorem 37 to our attention [BDEW24].

Gromov [G93] 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 37 that this is true for $X = K_4$.

Theorem 14.1.1. *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 37, 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 37. 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 [B65] 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

[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 14.2.1 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 14.5.8.

[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].

following, we will refer to the induced subgraphs of graphs in the Burling sequence as *Burling graphs*. From Theorem 14.1.1 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 14.1.2. *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 14.1.1, 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. □

14.1.4 An open conjecture

Even though Conjecture 36 is disproved for general graphs X [DHIM24], we would like to draw the reader's attention to the following special case[4] of Conjecture 36, which is still open. It is a coarse variant of Kuratowski's theorem and was first conjectured by Georgakopoulos and Papasoglu [GP25]:

Conjecture 39. *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.*

14.1.5 How this chapter is organised

In Section 14.2 we recall some important definitions and facts about the interplay between quasi-isometries and graph-decompositions of small radial width. Section 14.3 then contains some preparatory lemmas, which simplify the proofs of Theorems 37 and 38. Finally, we prove Theorem 38 in Section 14.4 and Theorem 37 in Section 14.5.

14.2 Preliminaries

We warn the reader that graphs in this chapter may be infinite, unless specified otherwise. Our notions mainly follow [D25]. 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 36 with any finite set of finite graphs.

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

14.2.1 Distance, radius and balls

Let G be a graph. We write $d_G(v,u)$ for the distance of the two vertices v and u in G. For two sets U and U' of vertices of G, we write $d_G(U,U')$ 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 $d_G(v,U') := d_G(\{v\},U')$ 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 $d_G(w,v) \leqslant k$ for every vertex v of G. If G is empty, then we define its radius to be 0. We remark that if G is disconnected but not the empty graph, then its radius is ∞. Note that G has radius at most k if and only if there is some vertex v of G with $V(G) = B_G(v,k)$. Additionally, if $U \subseteq V(G)$, then the *radius of U in G*, denoted as $\mathrm{rad}_G(U)$ is the smallest number $k \in \mathbb{N}$ such that there exists some vertex v of G with $U \subseteq B_G(v,k)$ or ∞ if such a $k \in \mathbb{N}$ does not exist.

If Y is a subgraph of G, then we abbreviate $d_G(U,V(Y))$, $\mathrm{rad}_G(V(Y))$, $B_G(V(Y),r)$, and $G[V(Y),r]$ as $d_G(U,Y)$, $\mathrm{rad}_G(Y)$, $B_G(Y,r)$, and $G[Y,r]$, respectively.

14.2.2 Fat minors

In this chapter we deviate from the definition of minor in [D25], 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 $d_G(Y,Z) \geqslant K$ 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.

14.2.3 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

[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 [D25].

(Q1) $M^{-1} \cdot d_H(h,h') - A \leqslant d_G(\varphi(h),\varphi(h')) \leqslant M \cdot d_H(h,h') + A$ for every $h,h' \in V(H)$, and

(Q2) for every vertex v of G, there exists a node h of H with $d_G(v,\varphi(h)) \leqslant A$.

We say that H is (M,A)-*quasi-isometric* to G if there exists a (M,A)-quasi-isometry from H to G. The following is a well-known fact.

Lemma 14.2.1. *If a graph H is (M,A)-quasi-isometric to a graph G, then G is $(M,3AM)$-quasi-isometric to H.* □

14.2.4 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* [DJKK26], 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 3.2], 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

$$\operatorname{oradw}_G(\mathcal{H}) := \sup_{h\in H} \operatorname{rad}_G(V_h),$$

and the *(inner-)radial spread* of $(H,\mathcal{V})$ is

$$\operatorname{irads}(\mathcal{H}) := \sup_{v\in Y} \operatorname{rad}_{H_v}(H_v).$$

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, § 3.3]. 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 14.2.2. *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 $C \in \mathcal{C}(G)$ such that $h \in V(H^C)$.*

In particular, we have that $\operatorname{oradw}_G(\mathcal{H}) = \sup_{C\in\mathcal{C}(G)} \operatorname{oradw}_C(\mathcal{H}^C)$ *and* $\operatorname{irads}(\mathcal{H}) = \sup_{C\in\mathcal{C}(G)} \operatorname{irads}(\mathcal{H}^C)$*. Moreover, if, for some connected graph X, we have $H^C \in \operatorname{Forb}_{\preccurlyeq}(X)$ for all $C \in \mathcal{C}(G)$, then also $H \in \operatorname{Forb}_{\preccurlyeq}(X)$.* □

14.2.5 Interplay between quasi-isometries and graph-decompositions

By [ADEFJKW23, § 3.3], 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 14.2.3. *[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 14.2.4. *[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 $(A(M^2+2) + M\lceil M(A+1)/2\rceil, 6MA + M + 1)$-radial H-decomposition.*

By Lemmas 14.2.3 and 14.2.4, Conjecture 36 can thus be reformulated in terms of graph-decompositions in the following qualitatively equivalent way.

Conjecture 36′. *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.

14.3 Preparatory work: sequences of graph-decompositions

As we have mentioned already in the introduction, we will prove Theorems 37 and 38 by showing the equivalent statements about graph-decompositions, that is, we will prove Conjecture 36′ 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 14.3.5) 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 14.3.1. *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 $C \in \mathcal{C}(G - Y^d)$ a partial graph-decomposition $\mathcal{H}^C = (H^C, \mathcal{V}^C)$ of G with support Y^C such that $N_G(C) \subseteq V(Y^C) \subseteq B_G(C, 1)$.

Let $\mathcal{H} = (H, \mathcal{V})$ be a partial graph-decomposition of G with support Y. For $C \in \mathcal{C}(G - Y)$, a partial graph-decomposition $\mathcal{H}^C = (H^C, \mathcal{V}^C)$ of G with support $Y^C \subseteq G[C, 1]$ is *feasible for $\mathcal{H}$ and C*, if $\partial_G C \subseteq V(Y^C)$ and there exist nodes h^C of H^C and g^C of H with $N_G(C) = V_{h^C}^C$ and $N_G(C) \subseteq V_{g^C}$. A family $\{\mathcal{H}^C \mid C \in \mathcal{C}(G - Y)\}$ of partial graph-decompositions of G is *feasible for $\mathcal{H}$* if each $\mathcal{H}^C$ is feasible for $\mathcal{H}$ and C.

Construction 14.3.2. Assume that $\mathcal{H} = (H, \mathcal{V})$ is a partial graph-decomposition of a graph G with support Y. Let $\{\mathcal{H}^C \mid C \in \mathcal{C}(G - Y)\}$ be a family of partial graph-decompositions $\mathcal{H}^C = (H^C, \mathcal{V}^C)$ of G with support Y^C which is feasible for $\mathcal{H}$.

We define the pair $\mathcal{H}' := (H', \mathcal{V}')$ as follows. Let H' be the graph obtained from the disjoint union of H and the H^C by identifying $h^C \in V(H^C)$ with $g^C \in V(H)$ for every $C \in \mathcal{C}(G - Y)$. For each node h of H', we set $V'_h := V_h$ if $h \in V(H)$ and $V'_h := V_h^C$ if $h \in V(H^C) \setminus \{h^C\}$ for some $C \in \mathcal{C}(G - Y)$. Set $Y' := Y \cup \bigcup_{C\in\mathcal{C}(G-Y)} Y^C$.

Note that the graph H' from Construction 14.3.2 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[6].

The following observation about $\mathcal{H}'$ from Construction 14.3.2 is immediate from its construction.

Observation 14.3.3. *Let $G, \mathcal{H}$ and the $\mathcal{H}^C$ be as in Construction 14.3.2. Then $\mathcal{H}'$ constructed as in Construction 14.3.2 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 $C \in \mathcal{C}(G - Y)$ and some 2-connected graph X, then $H' \in \mathrm{Forb}_{\preccurlyeq}(X)$,*

(ii) *if $\mathcal{H}$ and $\mathcal{H}^C$ are honest for every $C \in \mathcal{C}(G - Y)$, then $\mathcal{H}'$ is honest,*

(iii) $\mathrm{rad}_G(V'_h) = \mathrm{rad}_G(V_h)$, *if $h \in V(H)$,*

(iv) $\mathrm{rad}_G(V'_h) = \mathrm{rad}_G(V^C_h)$, *if $h \in V(H^C) \setminus \{h^C\}$ for some $C \in \mathcal{C}(G - Y)$,*

(v) $\mathrm{irads}(\mathcal{H}', v) = \mathrm{irads}(\mathcal{H}, v)$ *for $v \in V(Y) \setminus \partial_G Y$,*

(vi) $\mathrm{irads}(\mathcal{H}', v) = \mathrm{irads}(\mathcal{H}^C, v)$ *for $v \in V(Y^C) \setminus \partial_G Y = V(Y^C \cap C)$ for $C \in \mathcal{C}(G - Y)$, and*

(vii) $\mathrm{irads}(\mathcal{H}', v) \leqslant \mathrm{irads}(\mathcal{H}, v) + 2\sup_{C \in \mathcal{C}(G-Y)} \mathrm{irads}(\mathcal{H}^C, v)$ *for $v \in \partial_G Y$.* □

A partial graph-decomposition $\mathcal{H}$ of G with support Y is *R-component-feasible* for $R \in \mathbb{N}$ if for every component C of $G - Y$ there is $h \in V(H)$ with $N_G(C) \subseteq V_h$ and $\mathrm{rad}_G(V_h) \leqslant R$. A subgraph Y of G is *R-ball-componental* if for every component C of $G - Y$ there exists a vertex w of G such that C is a component of $G - B_G(w, R)$.

The following lemma shows that we may assume without loss of generality that every R-component-feasible partial graph-decomposition has an R-ball-componental support.

Lemma 14.3.4. *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* $\mathrm{oradw}(\mathcal{H}') \leqslant \max\{\mathrm{oradw}(\mathcal{H}), R\}$, $\mathrm{irads}(\mathcal{H}', v) = \mathrm{irads}(\mathcal{H}, v)$ *for $v \in V(Y)$, and* $\mathrm{irads}(\mathcal{H}', v) = 0$ *for $v \in V(Y') \setminus V(Y)$.*

Proof. For each node h of H with $\mathrm{rad}_G(V_h) \leqslant R$, we fix a vertex v_h of G such that $V_h \subseteq B_G(v_h, R)$. Since $(H, \mathcal{V})$ is R-component-feasible, we may further fix for each $C \in \mathcal{C}(G - Y)$, a node h_C of H such that $N_G(C) \subseteq V_{h_C}$ and $\mathrm{rad}_G(V_{h_C}) \leqslant R$. Set $U := V(Y) \cup \bigcup_{C \in \mathcal{C}(G-Y)}(B_G(v_{h_C}, R) \cap V(C))$. It is immediate from the construction that every component C' of $G - U$ is also a component of $G - B_G(v_{h_C}, R)$ where C is the

[6] 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 37 and 38 that the arising graph H' has no K_4^-- or K_4 minor, we will not do this here.

unique component of $G - Y$ with $C' \subseteq C$. Hence, the induced subgraph $Y' := G[U]$ is R-ball-componental.

It remains to extend $(H, \mathcal{V})$ to the desired graph-decomposition $(H, \mathcal{V}')$ of Y'. Set $V'_h := V_h \cup (B_G(v_h, R) \cap (\bigcup_{C \in \mathcal{C}(G-Y): h_C = h} V(C)))$ for every $h \in V(H)$. Note that, for every $h \in V(H)$ every vertex v of Y is contained in V'_h if and only if it is contained in V_h, and every vertex $v \in U \setminus V(Y)$ is contained in exactly one V'_h. Thus, $(H, \mathcal{V}')$ is indeed a graph-decomposition of Y', and its outer-radial width and radial spread are as claimed. Moreover, by construction, $N_G(C') \subseteq V'_{h_C}$ where C is the unique component of $G - Y$ with $C' \subseteq C$. Hence, $(H, \mathcal{V}')$ is R-component-feasible. □

Recall that our overall goal is to construct, for a given graph X and a graph G without K-fat X minor, a graph-decomposition of G along a graph H without X minor. For this, we recursively construct $f(K)$-radial partial decompositions $(H^d, \mathcal{V}^d)$ with support Y^d, where each $(H^d, \mathcal{V}^d)$ extends $(H^{d-1}, \mathcal{V}^{d-1})$ 'into the direction' of each component. Combining Observation 14.3.1 and Lemma 14.3.4, 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 14.3.5, 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 $H^C \in \mathrm{Forb}_{\preccurlyeq}(X)$ such that $N_G(C) \cup \partial_G C \subseteq V(Y^C) \subseteq B_G(C, 1)$, the neighbourhood $N_G(C)$ is a bag of $\mathcal{H}^C$, and $\mathrm{irads}(\mathcal{H}^C, v) \leqslant f''_1(K)$ for all $v \in N_G(C)$.

Theorem 14.3.5. *If a 2-connected finite graph X has property $(*)$ (with respect to functions R, f'_0, f'_1, f''_1), then Conjecture 36′ holds for X (with respect to $f := (f_0, f_1) := (f'_0, f'_1 + 2 \cdot f''_1 + 1)$).*

A partial graph-decomposition $\mathcal{H}$ of G with support Y is (R_0, R_1, R'_1)-*radial* in G if it is (R_0, R_1)-radial and $\mathrm{irads}(\mathcal{H}, v) \leqslant R'_1$ for every $v \in \partial_G Y$.

Proof. For $K \in \mathbb{N}$, let $R(K), f'_0(K), f'_1(K), f''_1(K)$ be given by property $(*)$ of X, and let G be any graph with no K-fat X minor. By Observation 14.2.2, 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 $V^0_h := V(Y^0)$. Clearly, $(H^0, \mathcal{V}^0)$ is honest, $R(K)$-component-feasible and $(f_0(K), 0, 0)$-radial, and $H^0 \in \mathrm{Forb}_{\preccurlyeq}(X)$.

Now let $d \in \mathbb{N}$ be given, and assume that for $i \leqslant d$ we have already constructed a sequence of honest, $R(K)$-component-feasible, $(f_0(K), f_1(K), f'_1(K))$-radial partial graph-decompositions $\mathcal{H}^i := (H^i, \mathcal{V}^i)$ of G with $R(K)$-ball-componental support Y^i and with $H^i \in \mathrm{Forb}_{\preccurlyeq}(X)$ such that each $\mathcal{H}^{i+1}$ extends $\mathcal{H}^i$. Let further $(H^C, \mathcal{V}^C)$, for every $C \in \mathcal{C}(G - Y^d)$, be given by property $(*)$ of X.

Then applying Construction 14.3.2 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 14.3.3, 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 14.3.4, $(H', \mathcal{V}')$ 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 14.3.1, 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. □

14.4 Proof of Theorem 38: Conjecture 36 for $X = K_4^-$

In this section we prove Theorem 38, which we restate here in the terminology of graph-decompositions.

Theorem 38′. *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 14.2.1 and 14.2.3, Theorem 38′ immediately implies Theorem 38.

To prove Theorem 38′, it suffices by Theorem 14.3.5 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 14.4.1 and Figure 14.1). It follows that the boundary of every component C of $G - B$ can be partitioned into two sets N_1, N_2 of small radius. If N_1 and N_2 are close in G, then the boundary of C has small radius, and we may obtain the desired partial graph-decomposition $\mathcal{H}^C$ of G by letting H^C be a K_2 on vertices h, h' and setting $V_h^C := N_G(C), V_{h'}^C := N_G(C) \cup \partial_G C$. Otherwise, the sets N_1 and N_2 are far apart in G. We then show that there is a partial path-decomposition $\mathcal{H}^C$ with support $Y^C \subseteq G[C, 1]$ as desired 'connecting' N_1 and N_2 such that its first and last bag are its only bags which intersect B (see Lemma 14.4.3 and Figure 14.2).

Lemma 14.4.1. *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 $d_G(u_i, u_j) \leqslant 2r$ for $i \neq j$ since B has radius r, and thus

$$r \geqslant d_G(u_i, u_j)/2 \geqslant 7K/2 > 3K.$$

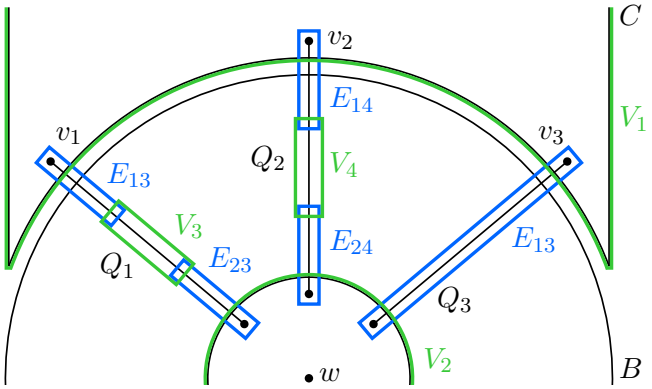

Figure 14.1: The K-fat K_4^- minor in the proof of Lemma 14.4.1.

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 14.1). Indeed, the sets V_i and the paths E_{ij} obviously form a model of K_4^-, so it remains to show that this K_4^--model is K-fat.

Since $G[V_1] = C$ is a component of $G - B_G(w, r)$ and $V_2 = B_G(w, r-(3K-1))$, we have $d_G(V_1, V_2) \geqslant 3K \geqslant K$. In particular, as the Q^i are shortest paths, we find $d_G(V_i, E_{(3-i)j}) \geqslant 2K \geqslant K$ as well as $d_G(E_{1j}, E_{2j}) \geqslant K$ for $i = 1, 2$ and $j = 3, 4$. By the choice of V_3, V_4 and the E_{ij} as subsets of the Q^i, it remains to show $d_G(Q^i, Q^j) \geqslant K$ for $i \neq j$. Indeed, we find

$$d_G(Q^i, Q^j) \geqslant d_G(u_i, u_j) - ||Q_i|| - ||Q_j|| \geqslant 7K - 2 \cdot 3K = K,$$

where the first inequality holds because any Q^i–Q^j path extends via Q^i and Q^j to a u_i–u_j path. □

For the next lemma, we need the following result, which is immediate from [ADEFJKW23, Proof of Lemma 4.2].

Lemma 14.4.2. *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 $2r+1$. Moreover, $V_p \subseteq B_G(p, 2r+1)$ for every $p \in V(P)$.

Lemma 14.4.3. *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 $K \in \mathbb{N}_{\geqslant 1}$.*

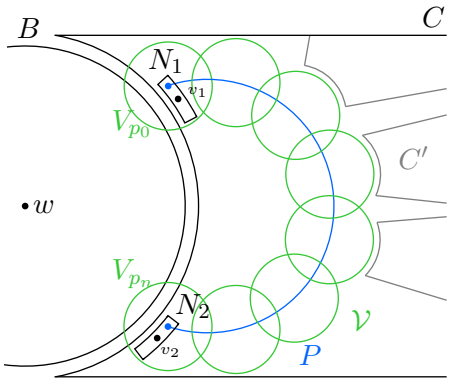

Figure 14.2: Depicted in green is the path-decomposition $(P, \mathcal{V})$ from Lemma 14.4.3 along the path P which is depicted in blue. By (ii), the neighbourhood of every component C' of $C - Y$ is contained in some bag in $\mathcal{V}$.

If G has no K-fat K_4^- minor, then there is an induced subgraph Y of $G[C, 1]$ which admits an honest $(28K+1, 28K+1)$-radial decomposition $(P, \mathcal{V})$ modelled on some path $P = p_0 \dots p_n$ such that

(i) *$N_G(C) \cup \partial_G C \subseteq V_{p_0} \cup V_{p_n}$, and $N_G(C)$ is disjoint from every V_p with $p \neq p_0, p_n$, and*

(ii) *for every component C' of $G - Y$ which meets, or equivalently is contained in, C, its neighbourhood $N_G(C')$ is contained in some bag V_p of $(P, \mathcal{V})$.*

Proof. Let us assume that G has no K-fat K_4^- minor and show that G then has a partial path-decomposition as claimed. Let v_1, v_2 be two vertices in $\partial_G C$ which are at least $42K+1$ apart. Then the balls B_1 and B_2 in G of radius $21K - 1 \leqslant \lfloor d_G(v_1, v_2)/2 \rfloor - 1$ around v_1 and v_2, respectively, are disjoint and also there is no B_1–B_2 edge in G. In particular, the sets $N_i := B_G(v_i, 7K - 1) \cap \partial_G C \subseteq B_i$ for $i \in [2]$ are at least $28K + 3$ apart because for all $n_1 \in N_1, n_2 \in N_2$ we have

$$d_G(n_1, n_2) \geqslant d_G(v_1, v_2) - d_G(n_1, v_1) - d_G(v_2, n_2) \geqslant (42K + 1) - 2 \cdot (7K - 1) = 28K + 3. \tag{14.1}$$

Further, since G has no K-fat K_4^- minor, Lemma 14.4.1 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 $N_{3-i} \subseteq V(C)$ 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

$$B_G(P, r_P) \subseteq B_G(C, r_P) = B_G(N_1, r_P) \cup V(C) \cup B_G(N_2, r_P) \subseteq B_1 \cup V(C) \cup B_2 = V(G'), \tag{14.2}$$

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,

$$B_{G'}(p, r_P) = B_G(p, r_P) \tag{14.3}$$

for all $p \in P$. So by Lemma 14.4.2, 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, $d_G(p, v) \leqslant 28K + 1$ for all $p \in V(P)$ and $v \in V'_p$.

Now set $Y := G[P, r_P] \cap G[C, 1]$. As we will see in a moment, the restriction of $(P, \mathcal{V}')$ to Y has already all desired properties, except that it does not satisfy the second part of (i). To solve this problem, we adapt $(P, \mathcal{V}')$ by deleting $N_G(C)$ from all bags V'_p of internal nodes $p \neq p_0, p_n$ of P. Formally, define the pair $(P, \mathcal{V})$ by $V_p := (V'_p \cap V(Y)) \setminus N_G(C)$ for $p \neq p_0, p_n$ and $V_p := V'_p \cap V(Y)$ for $p = p_0, p_n$. We claim that $(P, \mathcal{V})$ and Y are as desired.

We first prove that $(P, \mathcal{V})$ is indeed a graph-decomposition of Y. Since $(P, \mathcal{V}')$ is a graph-decomposition of $G[P, r_P]$, its restriction to $Y \subseteq G[P, r_P]$ is a graph-decomposition of Y. Thus, it suffices to show that every vertex $v \in N_G(C)$ is contained in precisely one of the bags V'_{p_0}, V'_{p_n} and that this bag also contains all neighbours of v.

Since $p_0 \in N_1$, we have $N_1 \subseteq B_G(p_0, 7K - 1)$. Otherwise, there is some $u \in N_1$ with $d_G(u, p_0) \geqslant 7K$. Since also $d_G(u, v_2), d_G(p_0, v_2) \geqslant d_G(N_1, N_2) \geqslant 7K$ by (14.1), Lemma 14.4.1 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 $d_G(p_0, N_2), d_G(p_n, N_1) \geqslant 28K + 3$ by (14.1) but $d_G(p, v) \leqslant 28K + 1$ for all $p \in V(P)$ and $v \in V'_p$, the set $B_G(N_1, 1)$ is disjoint from V'_{p_n}, and $B_G(N_2, 1)$ is disjoint from V'_{p_0}.

Further, since $(P, \mathcal{V}')$ has outer-radial width and radial spread at most $28K + 1$, so does $(P, \mathcal{V})$. It remains to show that $(P, \mathcal{V})$ satisfies (ii). Suppose for a contradiction that (ii) does not hold, that is, there is a component C' of $G - Y$ which is contained in C and which contains $z_1 \in N_G(B_G(p_{j_1}, r_P))$ and $z_2 \in N_G(B_G(p_{j_2}, r_P))$ with $0 \leqslant j_1 \leqslant j_2 \leqslant n$ such that $j_2 - j_1 > 2 \cdot (r_P + 1)$.

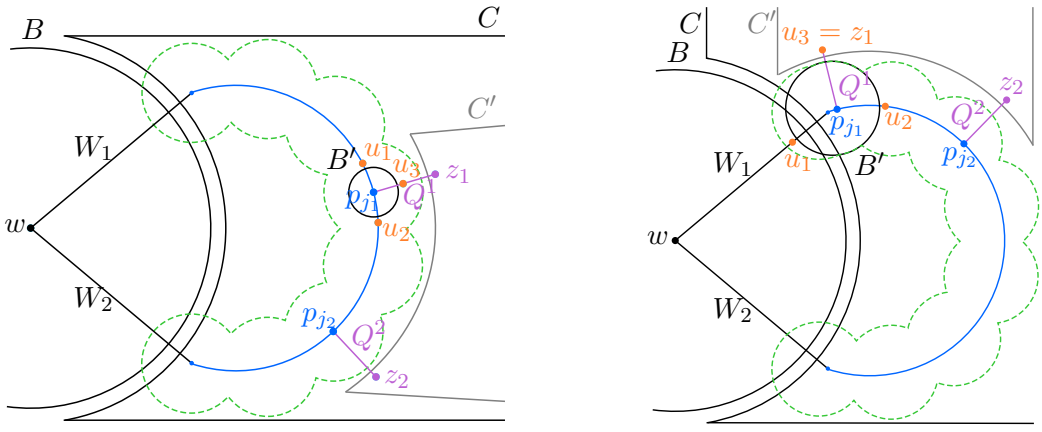

(a) Case 1: $j_1 \geqslant 7K$ and thus $u_1 \in V(P)$ (b) Case 2: $j_1 < 7K$ and thus $u_1 \in V(W_1)$

Figure 14.3: The situation in the proof of Lemma 14.4.3 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 $d_G(v_1, v_2)/2 \geqslant 14K + 1$.

Now first assume that $j_1 \geqslant 7K$ (see Figure 14.3a). 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 (14.3) 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 14.4.1 yields a K-fat K_4^- minor in G – a contradiction. Indeed, $Z := p_0Pu_1 \cup W^1 \cup W^2 \cup u_2Pp_n \cup Q^2 \cup C' \cup u_3Q^1z_1$ 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 14.3b). 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 14.4.1, 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 $d_G(p_0, p_n) > 2 \cdot (14K + 1)$ by (14.1) to show that Z is also disjoint from B'. □

We are now in a position to prove Theorem 38′. For this, by Theorem 14.3.5, it suffices to show that K_4^- has property (∗).

Lemma 14.4.4. *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 14.4.3 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 38′. By Theorem 14.3.5, Lemma 14.4.4 immediately yields Theorem 38′. □

Proof of Theorem 38. Let G be a graph with no K-fat K_4^- minor. By Lemma 14.2.3, Theorem 38′ 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 14.2.1 that G is $(84K+2, 3\cdot(84K+2)^2)$-quasi-isometric to H. Thus, Theorem 38 holds with $f(K) := (84K+2, 3\cdot(84K+2)^2)$. □

14.5 Proof of Theorem 37: Conjecture 36 for $X = K_4$

In this section we prove Theorem 37, which we restate here in the terminology of graph-decompositions.

Theorem 37′. *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 14.2.1 and 14.2.3, Theorem 37′ immediately implies Theorem 37.

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 $H + h_1h_2 \in \mathrm{Forb}_{\preccurlyeq}(K_4)$.[7] 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 14.5.1. *$\mathcal{H}_{SP}$ 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

$$\begin{aligned} R_0(K) &:= 129\cdot 5K + 5K = 130\cdot 5K,\\ R_0'(K) &:= 3R_0(K) + 5K + 1,\\ \ell(K) &:= 2R_0(K) + 5K + 2,\\ R_1(K) &:= 4\cdot(2(\ell(K) + 22K + 1) + 11K + 2), \text{ and}\\ R_2(K) &:= 2R_1(K) + 5K + 3. \end{aligned}$$

[7] 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 $H \in \mathrm{Forb}_{\preccurlyeq}(K_4)$.

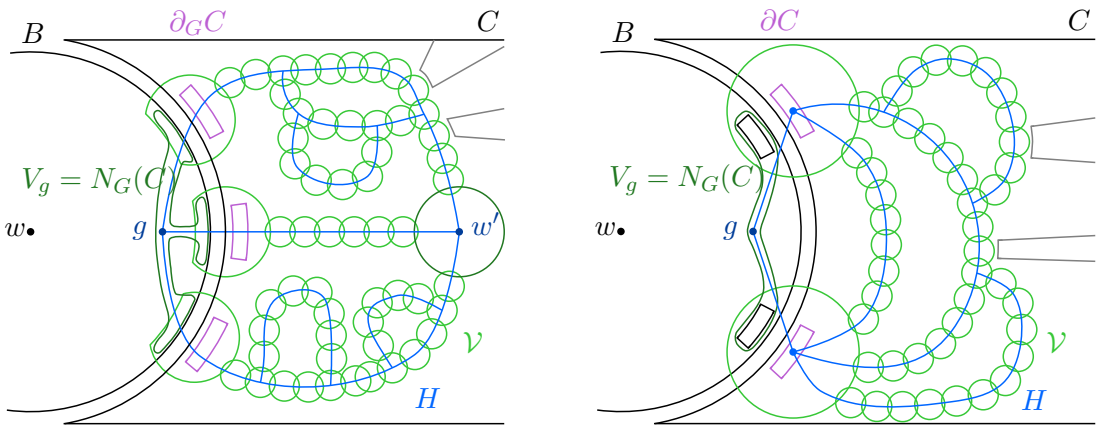

(a) Lemma 14.5.2 (b) Lemma 14.5.3

Figure 14.4: Depicted are the graph-decompositions $(H, \mathcal{V})$ with support Y in Lemmas 14.5.2 and 14.5.3 if G has no K-fat K_4 minor. By Lemma 14.5.2 (iv) and Lemma 14.5.3 (iv) the neighbourhood of every component of $C - Y$ is contained in some bag in $\mathcal{V}$.

Let us briefly sketch the proof of Theorem 37′. For this, let $K \in \mathbb{N}_{\geqslant 1}$, and let G be a graph with no K-fat K_4 minor. By Theorem 14.3.5 (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 14.5.4). So let such a ball B and component C of $G - B$ be given. To define $\mathcal{H}^C$, we distinguish two cases. If the boundary $\partial_G D$ of some component D of $C - B_G(B, 22K + 1)$ contains three vertices that are pairwise far apart, then we obtain the desired partial decomposition with support $Y^C \subseteq G[C, 1]$ by applying the following Lemma 14.5.2 to C and $C^* := D$.

Lemma 14.5.2. *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 14.5.3 (applied to $C := D$), and we can then identify the node of H^D given by Lemma 14.5.3 (i) with g'. This then yields the desired partial decomposition $\mathcal{H}^C$.

Lemma 14.5.3. *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 37′ from Lemmas 14.5.2 and 14.5.3. For this, as noted earlier, it suffices to prove that K_4 has property $(*)$.

Lemma 14.5.4. *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 $\mathcal{D}$ be the set of components of $C - B_G(B, 22K+1)$, and first assume that there is some $D \in \mathcal{D}$ such that $\partial_G D$ contains three vertices that are pairwise at least $R_1(K)$ apart. Then we obtain an induced subgraph $Y \subseteq G[C,1]$ with $\partial_G C \subseteq V(Y)$ and the desired honest $R_2(K)$-component-feasible $(f_0'(K), 7)$-partial decomposition $(H, \mathcal{V})$ of Y with $N_G(C) = V_g$ for some $g \in V(H)$ by applying Lemma 14.5.2 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 $\mathcal{D}' \subseteq \mathcal{D}$ be the set of components D of $C - B_G(B, 22K+1)$ such that there are two vertices in $\partial_G D$ which are at least $2R_1(K) + 5K + 2$ apart. Then by assumption, every $D \in \mathcal{D}'$ satisfies the premise of Lemma 14.5.3; let $(H^D, \mathcal{V}^D)$ be the partial decomposition with support Y^D obtained from applying Lemma 14.5.3 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 $D \in \mathcal{D} \setminus \mathcal{D}'$, let H^D be a K_2 on vertices g^D, h^D and set $V^D_{g^D} := V^D_{h^D} := N_G(D)$.

Now let H be the graph obtained from the disjoint union of the H^D for $D \in \mathcal{D}'$ by first identifying all g^D to a single node, which we call g', and adding a node g and the edge gg'. Note that $H \in \mathrm{Forb}_{\preccurlyeq}(K_4)$ since it is the 1-sum of the K_4-minor-free graphs H^D and K_2. Set $V_h := V^D_h$ for all $h \in V(H) \setminus \{g, g'\}$ where D is the unique component in $\mathcal{D}'$ such that $h \in V(H^D)$ and $V_{g'} := (B_G(B, 22K+1) \cap V(C)) \cup N_G(C)$ and $V_g := N_G(C)$. Then, by construction, $(H, \mathcal{V})$ is a partial decomposition with support $Y := G[\bigcup V(Y^D) \cup V_{g'}]$ since $V^D_{g^D} \subseteq V_{g'}$ for all $D \in \mathcal{D}'$.

We claim that $(H, \mathcal{V})$ is as desired. Indeed, the outer-radial width of $(H, \mathcal{V})$ is $R_2(K) + 22K + 1 \leqslant f'_0(K)$, since $\mathrm{rad}(V_{h^D}), \mathrm{rad}_G(V_g) \leqslant \mathrm{rad}_G(V_{g'}) \leqslant \mathrm{rad}_G(B) + 22K + 1 \leqslant R_2(K) + 22K + 1$ for every $D \in \mathcal{D} \setminus \mathcal{D}'$ and $\mathrm{rad}(V_h) \leqslant R_2(K) + 22K + 1$ for all other $h \in V(H)$ by (ii) of Lemma 14.5.3. Note that the radial spread of $(H, \mathcal{V})$ is $\leqslant 2 \cdot \max_{D \in \mathcal{D}} \mathrm{irads}((H^D, \mathcal{V}^D)) \leqslant 6 \leqslant f'_1(K)$, since $V_g \subseteq V_{g'}$, or it is 1 if $\mathcal{D} = \emptyset$.

By construction, it follows immediately from Lemma 14.5.3 (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 37′. By Theorem 14.3.5, Lemma 14.5.4 immediately yields Theorem 37′. □

Proof of Theorem 37. Let G be a graph with no K-fat K_4 minor. By Lemma 14.2.3, Theorem 37′ 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 14.2.1 that G is $(50470K + 142, 3 \cdot (50470K + 142)^2)$-quasi-isometric to H. Thus, Theorem 38 holds with $f(K) := (50470K + 142, 3 \cdot (50470K + 142)^2)$. □

The remainder of this section is devoted to the proofs of Lemmas 14.5.2 and 14.5.3. For this, we first show in Section 14.5.1 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 14.5.2 and 14.5.3, we prove Lemmas 14.5.2 and 14.5.3, respectively.

14.5.1 Finding a K-fat K_4 minor

Lemma 14.5.5. *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 $Q^i := q^i_0 \dots q^i_{n_i}$ be a shortest $B_G(v_1, r_1 - 2K)$–$B_G(v_2, r_2 - 2K)$ path in $G[B_1] \cup P_i \cup G[B_2]$; in particular, $n_i := ||P_i|| + 4K$.

Suppose that at least two $\mathring{P}_i$ lie in the same component C of $G - (B_1 \cup B_2)$. We show that G has a K-fat K_4 minor. For this, assume without loss of generality that $P_1 \subseteq G[C, 1]$.

Let W be a P_1–$B_G(Q^2 \cup Q^3, K)$ path in C; by symmetry, we may assume that W ends in $B_G(Q^2, K)$. Further, let $W' = w_0 \dots w_m$ be a shortest P_1–Q^2 path in $G[Q^2, K] \cup W$. Note that W' ends in a vertex in $q^2_K Q^2 q^2_{n_2-K} = Q^2 \cap G[P_2, K]$, as $W \subseteq C \subseteq G - (B_1 \cup B_2)$. In particular, since Q^2 is a shortest $B_G(v_1, r_1 - 2K)$–$B_G(v_2, r_2 - 2K)$ path in $G[B_1 \cup V(P_2) \cup B_2]$ it follows that $w_{m-K} W' w_m \subseteq G[P_2, 2K]$.

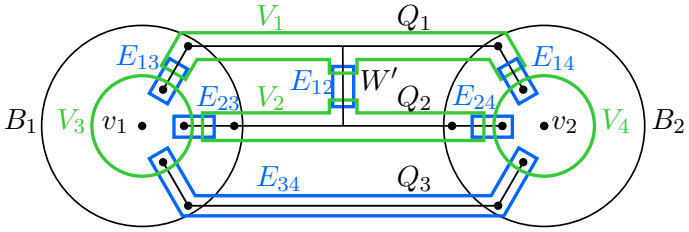

Figure 14.5: The K-fat K_4 minor in the proof of Lemma 14.5.5.

By defining the branch sets as

$$V_1 := V(q^1_K Q^1 q^1_{n_1-K}) \cup V(w_0 W' w_K),\ V_2 := V(q^2_K Q^2 q^2_{n_2-K} w_{m-K} W' w_m),$$
$$V_3 := B_G(v_1, r_1 - 2K) \text{ and } V_4 := B_G(v_2, r_2 - 2K).$$

and the branch paths as

$$E_{12} := w_K W' w_{m-K},\ E_{i3} := q^i_0 Q^i q^i_K,\ E_{i4} := q^i_{n_i-K} Q^i q^i_{n_i} \text{ and } E_{34} := Q^3$$

for $i \in [2]$ we obtain a K-fat K_4-model in G (see Figure 14.5). Indeed, the sets V_i and the paths E_{ij} obviously form a model of K_4. It remains to show that this K_4-model is K-fat.

By construction, we have $V_3, V(E_{i3}) \subseteq B_1$ and $V_4, V(E_{i4}) \subseteq B_2$ for $i \in [2]$, and thus by assumption

$$d_G(V_3, V_4), d_G(E_{i3}, E_{j4}), d_G(V_3, E_{i4}), d_G(V_4, E_{i3}) \geqslant d_G(B_1, B_2) \geqslant 5K \geqslant K$$

for $i, j \in [2]$. Moreover, $d_G(V(Q^1) \cup V_1, V(Q^2) \cup V_2) \geqslant d_G(P_1, P_2) - 2 \cdot 2K \geqslant 5K - 2 \cdot 2K = K$ since $V(Q^i) \cup V_i \subseteq B_G(P_i, 2K)$ for $i \in [2]$. This yields that

$$d_G(V_1, V_2), d_G(V_1, E_{2i}), d_G(V_2, E_{1i}), d_G(E_{1i}, E_{2i}) \geqslant K$$

for $i \in \{3, 4\}$. As above, we have $d_G(Q^3, V(Q^i) \cup V_i) \geqslant K$ since also $Q^3 \subseteq G[P_3, 2K]$ and $d_G(P_3, P_i) \geqslant 5K$ for $i \in [2]$. Further, $d_G(Q^3, W) \geqslant K$ by the choice of W. Thus,

$$d_G(E_{34}, V_i), d_G(E_{34}, E_{ij}), d_G(E_{34}, E_{12}) \geqslant K$$

for $i \in [2]$ and $j \in \{3, 4\}$. As C is a component of $G - (B_1 \cup B_2)$, we have $d_G(C, B_G(v_i, r_i - K)) = K + 1$ for $i \in [2]$. Since $E_{12} \subseteq W \subseteq C$ this implies that

$$d_G(E_{12}, E_{ij}), d_G(E_{12}, V_j) \geqslant d_G(C, B_G(v_i, r_i - K)) = K + 1 \geqslant K$$

for $i \in [2]$ and $j \in \{3, 4\}$. Finally, we obtain

$$d_G(V_i, V_j) \geqslant d_G(B_G(C, K), B_G(v_{j-2}, r_{j-2} - 2K)) = 2K - K = K$$

for $i \in [2]$ and $j \in \{3, 4\}$. This completes the proof. □

14.5.2 Proof of Lemma 14.5.3

Let us briefly sketch the proof of Lemma 14.5.3. 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 $d_G(B_1, B_2) \geqslant 5K + 2$. 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 14.5.8). To obtain the desired partial graph-decomposition $(H, \mathcal{V})$ to satisfy the conclusion of Lemma 14.5.3, 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 14.5.1 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 14.5.8, we construct the partial graph-decompositions $(H^D, \mathcal{V}^D)$ recursively via Lemma 14.5.7 below. For this, we will need the following coarse version of Menger's theorem for two paths.

Theorem 14.5.6. *[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.*[8]

Lemma 14.5.7. *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 $d_G(P, G[C, 1]) \geqslant 5K$. 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 14.6 for a depiction of the situation in Lemma 14.5.7.

Proof. Let Q be a shortest N_1–N_2 path through C. As $d_G(v_1, v_2) \geqslant r_1 + r_2 + 5K + 2$, the balls $B_G(B_1, \lceil (5/2)K \rceil)$ and $B_G(B_2, \lceil (5/2)K \rceil)$ are disjoint and not joined by an edge. Thus, Q is a shortest N_1–N_2 path in $G' := G[C, 1] \cup G[B_1 \cup B_2, \lceil (5/2)K \rceil]$. By applying Theorem 14.5.6 in G' to the sets N_1 and N_2 with $d = 5K$, we obtain either a vertex u of

[8] 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.

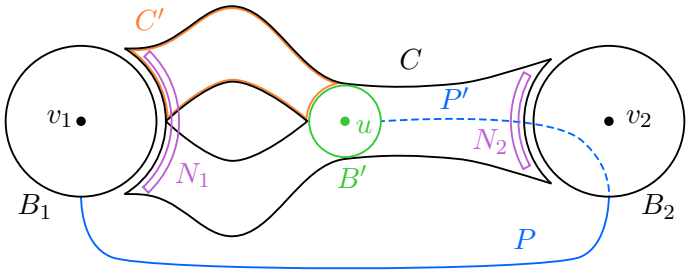

Figure 14.6: The situation in Lemma 14.5.7: A component C of $G - (B_1 \cup B_2)$ and the ball $B' \supseteq B$ around u. Depicted in blue is the path P and its dashed extension P', as required for the 'moreover' part applied to the orange component C'.

$Q \subseteq G[C, 1]$ such that every N_1–N_2 path in G' intersects $B_G(u, 129 \cdot 5K)$ or two N_1–N_2 paths Q_1 and Q_2 in G' such that $d_{G'}(Q_1, Q_2) \geqslant 5K$.

In the former case we are done (except for the 'moreover'-part) since u lies on a shortest N_1–N_2 path through C and $B' := B_G(u, 130 \cdot 5K) \supseteq B_G(u, 129 \cdot 5K)$; so suppose for a contradiction that the latter holds. Since $B_G(B_1, \lceil (5/2)K \rceil)$ and $B_G(B_2, \lceil (5/2)K \rceil)$ are disjoint and not joined by an edge, the internal vertices of any N_1–N_2 path in G' are in C; in particular, $\mathring{Q}_1, \mathring{Q}_2 \subseteq C$. Moreover, $d_G(Q_1, Q_2) \geqslant 5K$ because $d_{G'}(Q_1, Q_2) \geqslant 5K$ and any C-path in G (and hence any $\mathring{Q}_1$–$\mathring{Q}_2$ path) meeting $G - G'$ has length at least $2 \cdot (\lceil (5/2)K \rceil + 1) \geqslant 5K + 2$. Since also $d_G(Q_1 \cup Q_2, P) \geqslant d_G(G[C, 1], P) \geqslant 5K$ by assumption, the paths Q_1, Q_2, P are pairwise at least $5K$ apart. Thus, as $\mathring{Q}_1$ and $\mathring{Q}_2$ lie in the same component C of $G - (B_1 \cup B_2)$, applying Lemma 14.5.5 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 $d_G(P', G[C', 1]) \geqslant 5K$. For this, let S be a shortest $G[C', 1]$–P' path in G. We show that S has length at least $5K$. Then $d_G(G[C', 1], P') \geqslant ||S|| \geqslant 5K$, which yields the claim. Since $d_G(G[C', 1], P) \geqslant d_G(G[C, 1], P) \geqslant 5K$ as $C' \subseteq C$, we are done if S ends in P. Hence, we may assume that S ends in P_1 or P_2.

We distinguish two cases. First, assume that S meets $B_G(w, 129 \cdot 5K) \subseteq B'$. Since $P_1 \cup P_2$ does not meet $B' = B_G(B_G(w, 129 \cdot 5K), 5K)$ except in its last vertex, we have $d_G(P_1 \cup P_2, B_G(w, 129 \cdot 5K)) \geqslant 5K$. 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 $d_G(v_1, v_2) \geqslant 2r_1 + r_2 + 5K + 2$, we have $d_G(B_1, B_2) \geqslant 5K + 2$, and thus S has length at least $5K$. □

Lemma 14.5.8. *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 $d_G(P, G[C, 1]) \geqslant 5K$.*

If G has no K-fat K_4 minor, then there exists an induced subgraph Y of $G[C, 1]$ which admits an honest decomposition $(H, \mathcal{V})$ modelled on a graph $H \in \mathcal{H}_{SP}$ with terminals h_1, h_2 such that

(a) *$V_{h_i} = N_G(C) \cap B_i$ for $i \in [2]$,*

(b) *$\mathrm{rad}_G(V_h) \leqslant R'_0(K)$ for all $h \in V(H)$ with $h \notin B_H(\{h_1, h_2\}, 1)$,*

(c) *$d_G(v_i, v) \leqslant r_i + 2 \cdot R_0(K) + 5K + 1$ for $i \in [2]$ and all $v \in V_h$ with $h \in B_H(h_i, 1)$,*

(d) *$\mathrm{irads}(\mathcal{H}) \leqslant 3$; moreover, $d_H(h_i, h) \leqslant 3$ for all $h \in V(H_v)$ and $v \in N_G(C) \cap B_i$ for $i \in [2]$, and*

(e) *for every component C' of $G - Y$ which meets, or equivalently is contained in, C, its neighbourhood $N_G(C')$ is contained in some bag V_h of $(H, \mathcal{V})$.*

We remark that the proof of Lemma 14.5.8 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 14.5.8 for finite graphs. We may apply Lemma 14.5.7 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 14.5.8 for arbitrary graphs. We define the desired graph-decomposition $(H, \mathcal{V})$ recursively (see Figure 14.7 for a sketch). Initialise H^0 as a K_2 on vertices $h_1 := v_1$ and $h_2 := v_2$. In particular, $H^0 \in \mathcal{H}_{SP}$ with terminals h_1 and h_2. Set $V_{h_i} := N_G(C) \cap B_i$ and $r_{h_i} := r_i$ for $i \in [2]$, and note that $N_G(C) \subseteq V_{h_1} \cup V_{h_2}$. Set further $\mathcal{C}^0 := \{C\}$, $w_1^C := h_1$ and $w_2^C := h_2$. We warn the reader that in the following the vertices of the H^i are vertices of G, i.e. $V(H^i) \subseteq V(G)$. Hence, we may divert from our usual convention, and denote the vertices of H by u, w.

Now suppose that for $n \in \mathbb{N}$ we have constructed graphs $H^0 \subseteq \ldots \subseteq H^n \in \mathcal{H}_{SP}$ with terminals h_1 and h_2, and sets $V_u \subseteq B_G(C, 1)$ for all $u \in V(H^n)$. Assume further that, for every $i \leqslant n$, we fixed

(I) the collection $\mathcal{C}^i$ of all components D of $C - \bigcup_{u \in H^i} V_u$ whose neighbourhood is not contained in a V_u for some $u \in V(H^i)$,

(II) for every $D \in \mathcal{C}^i$, an edge $w_1^D w_2^D$ of H^i,

(III) for every $D \in \mathcal{C}^{i-1} \setminus \mathcal{C}^i$, some $u_D \in V(H^i) \setminus V(H^{i-1})$ such that $w_1^D u_D, u_D w_2^D \in E(H^i)$, if $i \geqslant 1$,

(IV) for every $u \in V(H^i) \setminus V(H^{i-1})$, some $C_u \in \mathcal{C}^{i-1}$ with $u \in B_G(C_u, 1)$ and $u_{C_u} = u$, if $i \geqslant 1$, and

(V) for all $u \in V(H^i) \setminus \{h_1, h_2\}$, the bag $V_u := B_G(u, R_0(K)) \cap B_G(C_u, 1)$ and $r_u := R_0(K)$,

such that

(i) for every edge e of $G[\bigcup_{u \in H^i} V_u]$ that is not an edge of $\bigcup_{u \in H^i} G[V_u]$ there is an edge $uu' \in E(H^i)$ such that $e \in G[V_u \cup V_{u'}]$;

(ii) every $D \in \mathcal{C}^i$ is a component of $G - (V_{w_1^D} \cup V_{w_2^D})$;

(iii) for every $u \in V(H^i) \setminus V(H^{i-1})$ every $V_{w_1^{C_u}}$–$V_{w_2^{C_u}}$ path through C_u intersects $B_G(u, 129 \cdot 5K)$;

(iv) if $i \geqslant 1$, then for every $D \in \mathcal{C}^{i-1}$ we have $D \notin \mathcal{C}^i$ if and only if $d_G(w_1^D, w_2^D) \geqslant r_{w_1^D} + r_{w_2^D} + 5K + 2$;

(v) for every $D \in \mathcal{C}^i \setminus \mathcal{C}^{i-1}$ and the unique component $D' \in \mathcal{C}^{i-1} \setminus \mathcal{C}^i$ with $D \subseteq D'$ we have $w_1^D = w_1^{D'}$, $w_2^D = u_{D'}$ or $w_1^D = u_{D'}$, $w_2^D = w_2^{D'}$;

(vi) for every $D \in \mathcal{C}^i$ there exists a $B_G(w_1^D, r_{w_1^D})$–$B_G(w_2^D, r_{w_2^D})$ path P which satisfies $d_G(P, G[D, 1]) \geqslant 5K$.

In the following, we will refer to components $D \in \mathcal{C}^n$ as in (iv), that is those which satisfy $d_G(w_1^D, w_2^D) \geqslant r_{w_1^D} + r_{w_2^D} + 5K + 2$, as *long*.

We now obtain H^{n+1} as follows. For every long component $D \in \mathcal{C}^n$, let $u_D \in B_G(D, 1)$ be a vertex obtained from the application of Lemma 14.5.7 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 $D \in \mathcal{C}^n$ the vertex u_D and the edges $w_1^D u_D$ and $u_D w_2^D$.

Let the collection $\mathcal{C}^{n+1}$, the sets V_{u_D} and the integer r_{u_D} be as required by (I) and (V), respectively. The construction obviously ensures (III) and (iv). For (IV), we set $C_{u_D} := D$ for every $u_D \in V(H^{n+1}) \setminus V(H^n)$, which also yields (iii). For (II), let $D' \in \mathcal{C}^{n+1} \setminus \mathcal{C}^n$. Then D' is contained in a long component $D \in \mathcal{C}^n$. As the neighbourhood of D' is not contained in a V_u for $u \in V(H^{n+1})$, the ball $B_G(u_D, 129 \cdot 5K)$ meets every $V_{w_1^D}$–$V_{w_2^D}$ path through D and $B_G(u_D, 129 \cdot 5K) \cap V(D) \subseteq V_{u_D}$, the component $D' \subseteq D$ attaches to V_{u_D} and precisely one of $V_{w_1^D}$ and $V_{w_2^D}$, say $V_{w_1^D}$. Then we set $w_1^{D'} := w_1^D$ and $w_2^{D'} := u_D$, which immediately ensures (ii) and (v). The moreover-part of Lemma 14.5.7 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 $D \in \mathcal{C}^n$ was already an edge of $G[\bigcup_{u \in H^n} V_u]$, and hence satisfies (i) by (i) of $(H^n, \mathcal{V}^n)$. Now let e be any edge of $G[\bigcup_{u \in H^{n+1}} V_u]$ which is not in $G[\bigcup_{u \in H^n} V_u]$. Then there is $D \in \mathcal{C}^n$ such that $e = xy$ has one endvertex, say x, in $V_{u_D} \cap V(D)$, so e has its other endvertex y in $G[D, 1]$. Hence, by (ii) of $(H^n, \mathcal{V}^n)$, the vertex y is in $V_{w_1^D} \cup V_{w_2^D}$. By construction, $w_1^D u_D, u_D w_2^D \in E(H^{n+1})$, and hence e satisfies (i). Also H^{n+1} is still in

$\mathcal{H}_{SP}$ with terminals h_1, h_2 by (II) and Proposition 14.5.1. This completes the verification that H^{n+1} is as desired.

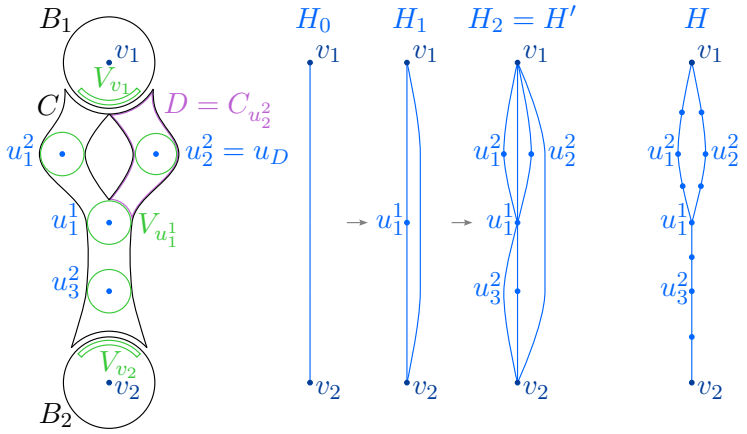

Figure 14.7: Depicted are the graphs H^0, H^1, H^2, H in blue and the bags V_u for the nodes u of H^2 for the component C of $G - (B_1 \cup B_2)$ in green. In this example, we have $w_1^D = v_1, w_2^D = u_1^1$ for the component $D = C_{u_2^2}$ in purple, and $H' = H^2$.

Let $H' := \bigcup_{i \in \mathbb{N}} H^i$ be the limit of the H^i. By Proposition 14.5.1, 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 $d_G(u, u') \geqslant r_u + r_{u'} + 5K + 2$ and then subdividing each remaining edge precisely once. By Proposition 14.5.1, $H \in \mathcal{H}_{SP}$ 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 $C_{h_v} \in \mathcal{C}^{n-1}$. Thus, the choice of h_v is unique by (IV) and (V). If $n = 0$, then we have $v \in N_G(C) \cap B_i$ for a unique $i \in [2]$ and $h_v = h_i$. We first prove that $H'_v = H'[\{h \in V(H') \mid v \in V_h\}]$ 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 $C_{h'} \in \mathcal{C}^m$. Since n is minimal and h_v is unique, we have that $m > n$ and that $C_{h'}$ attaches to v. For every ℓ with $m \geqslant \ell > n$, let C^ℓ be the unique component in $\mathcal{C}^\ell$ which contains $C_{h'}$; in particular, $C_{h'} = C^m$. Since $C_{h'} \in \mathcal{C}^m \setminus \mathcal{C}^{m-1}$, it follows from (I), (iv) and (V) that no C^ℓ is in $\mathcal{C}^{\ell-1}$. Also, every C^ℓ with $\ell > n$ attaches to v as $v \in V_{h_v}$ and $C^m = C_{h'}$ attaches to v. By iteratively combining this with (ii) and (v), we obtain that $h_v = w_2^{C^{n+1}} = \cdots = w_2^{C^m}$ (after possibly swapping the names of $w_1^{C^\ell}$ and $w_2^{C^\ell}$). Moreover, by (IV), $h' = u_{C_{h'}} = u_{C^m}$, and thus $h_v h' = w_2^{C^m} u_{C^m}$ is an edge of H' by (III), which completes the proof that $H'_v \subseteq H'[h_v, 1]$.

Then (H2) of $(H, \mathcal{V})$ follows: For every node h' of H' with $v \in V_{h'}$, since $d_G(h_v, h') \leqslant r_{h_v} + r_{h'}$ as V_{h_v} and $V_{h'}$ intersect in v, we did not delete the edge $h_v h'$ in the construction of H. Since H is obtained from H' by subdividing each remaining edge precisely once, and since we associated to every subdivision vertex u of H the union V_u of the two bags associated with its neighbours, $H'_v \subseteq H'[h_v, 1]$ implies that $H_v = H[\{h \in V(H) \mid v \in V_h\}]$

is contained in $H[h_v, 3]$.[9] In particular, H_v is connected, which yields (H2). Moreover, $\operatorname{irads}((H, \mathcal{V})) \leqslant 3$ 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 $d_G(v, u) \leqslant r_u$ 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, $C' \in \mathcal{C}^n$ for all such n, as its neighbourhood is not even contained in some V_h with $h \in V(H')$. Let m be the minimal such n. By (II) and (ii), C' is a component of $G - (V_g \cup V_{g'})$ and $gg' = w_1^{C'} w_2^{C'}$ is an edge of H^n. If we did not delete gg' when constructing H, then we have a contradiction, as the bag V_u associated with the subdivision vertex u of H on the edge gg' then contains the entire neighbourhood of C'. Otherwise, we have $d_G(g, g') > r_g + r_{g'} + 5K + 2$, which by (iv) contradicts that $C' \in \mathcal{C}^{m+1}$. □

Proof of Lemma 14.5.3. 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 $d_G(P, G[C, 1]) \geqslant R_1(K) - 1 \geqslant 5K$. Since also $d_G(v_1, v_2) \geqslant 2R_1(K) + 5K + 1$ by assumption, we may apply Lemma 14.5.8 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 14.5.8 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 $Y := G[(\bigcup V(Y^{C'})) \cup ((B_1 \cup B_2) \cap V(C)) \cup N_G(C)]$. The graph H is obtained from the disjoint union of the $H^{C'}$ and three new vertices h_1, h_2 and g by identifying all $h_1^{C'}$ with h_1 as well as all $h_2^{C'}$ with h_2 and adding the edges $h_1 g$ and $h_2 g$. Note that this is a parallel composition of graphs $H^{C'} \in \mathcal{H}_{SP}$ and a path of length 2; so H is again in $\mathcal{H}_{SP}$ with terminals h_1 and h_2 by Proposition 14.5.1; in particular, $H \in \operatorname{Forb}_{\preccurlyeq}(K_4)$. 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 $\partial_G C$

[9] We remark that H_v is not necessarily contained in $H[h_v, 2]$: for any given $h' \in H'[h_v, 1]$ and any edge $e \neq h'h_v$ incident with h', the bag corresponding to the subdivision vertex of e in H also contains v but has distance 3 to h_v.

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 $\operatorname{rad}(V_h) \leqslant R_1(K) + 2R_0(K) + 5K + 2$ 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 $\operatorname{rad}_G(V_h) \leqslant R_1(K) + 2R_0(K) + 5K + 2$, as we have shown for (ii) above. Hence, $(H, \mathcal{V})$ is as desired. □

14.5.3 Proof of Lemma 14.5.2

Let us briefly sketch the proof of Lemma 14.5.2. 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 14.5.10). 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 14.5.5. 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 14.5.8 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 14.5.11). 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 $G[C, 1]$.

We first show that there exists a vertex w' as described above.

Lemma 14.5.9. *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

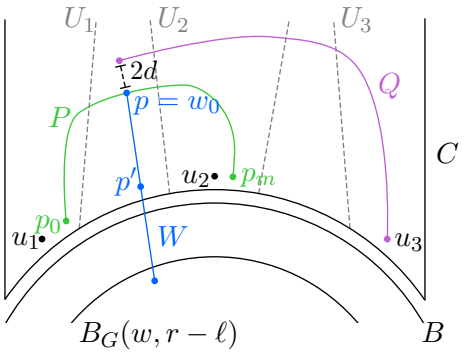

Figure 14.8: Setup in the beginning of the proof of Lemma 14.5.9

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 14.8). Since $p \notin U_3$, we may choose W so that $d_G(p', u_3) \geqslant 2\ell + d + 2$. Moreover, as $d_G(u_1, u_2) \geqslant 4 \cdot (2\ell + d + 2)$, we cannot have $d_G(p', u_1), d_G(p', u_2) < 2 \cdot (2\ell + d + 2)$; by symmetry, we may thus assume $d_G(p', u_2) \geqslant 2 \cdot (2\ell + d + 2)$. In particular, this implies that

$$d_G(p_m, p') \geqslant d_G(p', u_2) - d_G(u_2, p_m) \geqslant 2 \cdot (2\ell + d + 2) - (2\ell + d + 2) = 2\ell + d + 2. \quad (14.4)$$

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 14.9a).

For this, let Q_2 be a shortest p_m–$B_G(w, r-\ell)$ path and Q_3 a shortest u_3–$B_G(w, r-\ell)$ path. Further, let S_3 be a shortest x_3–w_{i_3} path and let s_3 be the first vertex on the path $S'_3 := x_3 S_3 w_{i_3} W w_{i_2}$ that is contained in $B_G(w_{i_2}, 2d)$ (see Figure 14.9a). We now define our final three $B_G(w, r-\ell)$–$B_G(w', 2d)$ paths P_1, P_2, P_3 (see Figure 14.9b):

$$P_1 := w_{i_2+2d} W w_k, \; P_2 := x_2 P p_m Q_2 \text{ and } P_3 := s_3 \bar{S}'_3 x_3 \bar{Q} u_3 Q_3.$$

We claim that the P_i are as desired, i.e. $d_G(P_i, P_j) \geqslant d$ for every two distinct $i, j \in [3]$. First, we show $d_G(P_1, P_2) \geqslant d$: By the choice of x_2, we have $d_G(W, x_2 P p_m) \geqslant 2d \geqslant d$. Further, since Q_2 starts in p_m and both shortest $\partial_G C$–$B_G(w, r-\ell)$ paths Q_2 and $p' W w_k$ have length $\ell + 1$, we have

$$d_G(p' W w_k, Q_2) \geqslant d_G(p_m, p') - ||p' W w_k|| - ||Q_2|| \geqslant (2\ell + d + 2) - 2(\ell + 1) = d. \quad (14.5)$$

Since $d_G(p_m, p') \geqslant 2\ell + d$ by (14.4), $\ell \geqslant d$, Q_2 is a shortest $\partial_G C$–$B_G(w, r-\ell)$ path and $p' \in \partial_G C$, we have $d_G(Q_2, B_G(p', d) \cap V(W)) \geqslant d$. As W is a shortest p–$B_G(w, r-\ell)$ path,

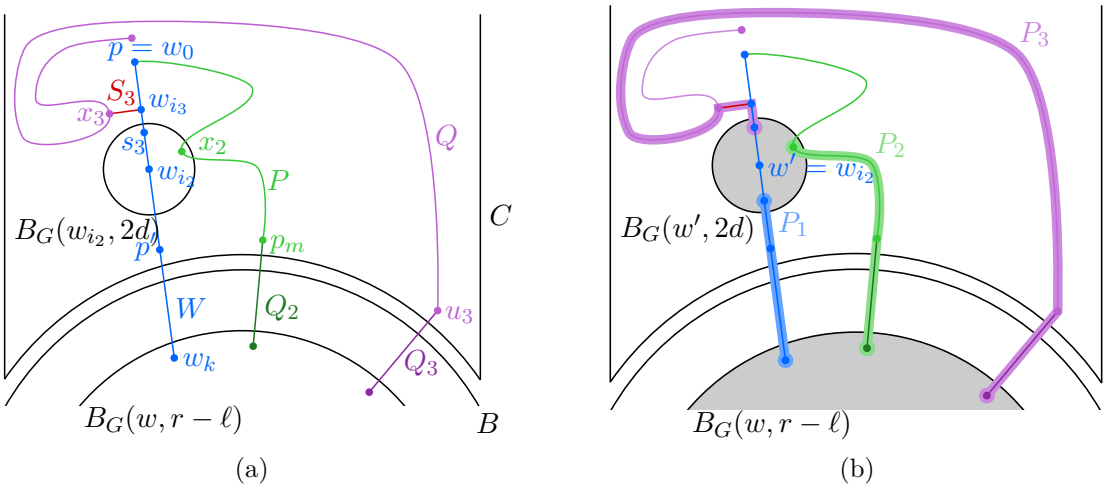

Figure 14.9: Choice of w' and construction of the paths P_1, P_2, P_3 in the proof of Lemma 14.5.9.

the remainder of W in C has distance $\geqslant d$ to B and hence to Q_2. All in all, $d_G(W, Q_2) \geqslant d$, and hence $d_G(P_1, P_2) \geqslant d$.

Next, we show $d_G(P_1, P_3)$: Analogously to $d_G(W, Q_2) \geqslant d$, we also find $d_G(W, Q_3) \geqslant d$. Moreover, $d_G(W, x_3\bar{Q}u_3) \geqslant 2d \geqslant d$ by the choice of x_3. Thus, to see that $d_G(P_1, P_3) \geqslant d$, it suffices to verify $d_G(w_{i_2+2d}Ww_k, S_3') \geqslant d$: Since W is a shortest path in G and $i_3 \leqslant i_2$, we have $d_G(w_{i_2+2d}Ww_k, w_{i_3}Ww_{i_2}) = d_G(w_{i_2+2d}, w_{i_2}) = 2d$. Moreover, by the choice of $x_3 \in V(Q)$ and S_3, the x_3–w_{i_3} path S_3 has length $||S_3|| = d_G(W, x_3) = 2d$. Thus, the distance from the first half of S_3 to $w_{i_2+2d}Ww_k$ is $\geqslant d_G(x_3, W) - ||S_3||/2 = 2d - d = d$, and the distance from the second half of S_3 to $w_{i_2+2d}Ww_k$ is $\geqslant d_G(w_{i_3}, w_{i_2+2d}Ww_k) - ||S_3||/2 \geqslant 2d - d = d$.

Finally, we show $d_G(P_2, P_3) \geqslant d$: By the choice of Q, we have $d_G(P, Q) \geqslant 2d$. Similarly to the previous argument, we find $d_G(x_2Pp_m, S_3') \geqslant d$, as $d_G(W, x_2Pp_m), d_G(P, Q) \geqslant 2d$, and thus the distance from the first half of S_3 to x_2Pp_m is $\geqslant d_G(x_3, P) - ||S_3||/2 \geqslant 2d - d = d$, and the distance from the second half of S_3 to x_2Pp_m is $\geqslant d_G(w_{i_3}, x_2Pp_m) - ||S_3||/2 \geqslant 2d - d = d$. Moreover, analogously to (14.5), we find

$$\begin{aligned} d_G(Q_2, Q_3) &\geqslant d_G(p_m, u_3) - ||Q_2|| - ||Q_3|| \\ &\geqslant d_G(U_2 \cap \partial_G C, u_3) - 2(\ell+1) \geqslant (2\ell + d + 2) - 2(\ell+1) = d. \end{aligned}$$

Further, $d_G(x_2Pp_m, Q_3), d_G(u_3Qx_3, Q_2) \geqslant d$ since $d_G(P, u_3), d_G(Q, p_m) \geqslant d_G(P, Q) \geqslant 2d$ and $P, Q \subseteq C \subseteq G - B$, while Q_2, Q_3 are shortest $\partial_G C$–$B_G(w, r-\ell)$ paths starting in p_m, u_3, respectively. It thus remains to show $d_G(S_3, Q_2) \geqslant d$: For this, it suffices to prove $d_G(x_3, Q_2) \geqslant 3d$, as S_3 has length $2d$ and starts in x_3. If $d_G(x_3, B) > 3d$, then $d_G(x_3, Q_2) \geqslant 3d$ follows immediately, as $Q_2 \subseteq G[B, 1]$. Thus, we may assume $d_G(x_3, B) \leqslant 3d$. Suppose for a contradiction that $d_G(x_3, Q_2) < 3d$. Then $d_G(x_3, p_m) \leqslant 2d_G(x_3, Q_2) < 6d$, since $x_3 \in V(C)$ and Q_2 is a shortest p_m–$B_G(w, r-\ell)$ path. Also $d_G(w_{i_3}, p') \leqslant d_G(w_{i_3}, x_3) + d_G(x_3, B) \leqslant 5d$, since W is a shortest path. Thus, $d_G(p_m, p') \leqslant d_G(p_m, x_3) + d_G(x_3, w_{i_3}) + d_G(w_{i_3}, p') < 13d$, which contradicts $d_G(p_m, p') \geqslant 2\ell + d$ as $\ell \geqslant 6d$. □

Corollary 14.5.10. *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 $d_G(P, G[D, 1]) \geqslant 5K$.

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 14.5.9 (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 $d_G(P_i, P_j) \geqslant 11K$ for all $i \neq j \in [3]$. Set $A := B_G(w', 22K)$, and let D be a component of $G - (B' \cup A)$ that attaches to B' and A.

We show that there is some $i \in [3]$ such that $d_G(P_i, G[D, 1]) \geqslant 5K$. Towards a contradiction, suppose that $d_G(P_i, D) \leqslant 5K$ for all $i \in [3]$. It then follows that all three paths P_i lie in the unique component D' of $G - (B_G(w, r - \ell - 3K) \cup B_G(w', 19K))$ that contains D. For each $i \in [3]$, let P_i' be a shortest $B_G(w, r - \ell - 3K)$–$B_G(w', 19K)$ path in $P_i \cup G[B'] \cup G[A]$. Then $d_G(P_i', P_j') \geqslant 11K - 2 \cdot 3K = 5K$, which by Lemma 14.5.5 ($d_G(w, w') > r_1 + r_2 + 5K + 1$ with $r_1 = r - \ell - 3K, r_2 = 19K$, since $d_G(w', w) \geqslant r$ and $\ell \geqslant 23K$) implies that G has a K-fat K_4 minor, which contradicts our assumption. □

Lemma 14.5.11. *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 $d_G(P, G[D, 1]) \geqslant 5K$.*

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 14.5.8 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 14.5.8.

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 $v \in V^1_h \cap B_2^{+1} \subseteq B_G(D, 1)$, 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 $H_{Q^1} := H[\{h \in V(H^1) \mid V^1_h \cap Q^1 \neq \emptyset\}]$ since $V^1_{h_2} = N_G(D) \cap B'_2$ 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 $V^1_{h_2}$ that contain a vertex of B_2^{+1}, i.e. $V^2_{h_2} = \bigcup_{h \in \tilde{H}} V^1_h$ and $V^2_h = V^1_h$ 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 $V_h := V^2_h = V^1_h$ for all $h \in V(H) \setminus \{h_1, h_2\}$ and $V_{h_i} := V^2_{h_i} \cup B_i$ 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, $V^2_{h_2}$ is the only bag of $(H, \mathcal{V}^2)$ that may contain vertices of B_2, and because $B_1 \setminus V^1_{h_1}$ does not meet Y^2 and B_2 as $Y^1 \subseteq G[D, 1]$ and because B_1 and B_2 are disjoint.

To see that (H1) holds, observe that, since $(H, \mathcal{V}^2)$ is a graph-decomposition of Y^1, and $V_{h_2} \supseteq B_2$ and $V_{h_1} = B_1$, it suffices to show that the parts of $(H, \mathcal{V})$ cover all edges between $Y^1 \cap D$ and $B_1 \cup B_2$ as well as all edges between B_1 and B_2. Since $N_G(D) \cap B_1 \subseteq V^2_{h_1}$ and $(H, \mathcal{V}^2)$ satisfies (H1), the parts of $(H, \mathcal{V}^2)$ already covered all edges between B_1 and $Y^1 \cap D$, and so the parts of $(H, \mathcal{V})$ do so as well. Further, the parts of $(H, \mathcal{V})$ cover all edges between B_2 and $Y^1 \cap D$ since $V(Y^1) \cap N_G(B_2) \subseteq V(Y^1) \cap B_2^{+1} \subseteq V_{h_2}$ as ensured by the contraction of $\tilde{H}$ in the construction of $(H^2, \mathcal{V}^2)$. Finally, there are no edges between B_1 and B_2 because $d_G(v_1, v_2) \geqslant r_1 + 2 + r_2$.

We now check that $(H, \mathcal{V})$ satisfies (α) – (ε). By definition, $V_{h_2} = V^2_{h_2} \cup B_2 \supseteq B_2$ and by (a) $V_{h_1} = V^2_{h_1} \cup B_1 = B_1$, and hence $(H, \mathcal{V})$ satisfies (α). To see that $(H, \mathcal{V})$ satisfies (β), let us first note that $\operatorname{rad}(V_{h_1}) = \operatorname{rad}_G(B_1) \leqslant r_1 \leqslant R_0(K)$.

So by (b) and (c), we have $\operatorname{rad}_G(V_h) = \operatorname{rad}_G(V^1_h) \leqslant \max\{r_1 + 2R_0(K) + 5K + 1, R'_0(K)\} \leqslant R'_0(K)$ for all nodes $h \neq h_1$ of $V(H^1)$ with $h \notin B_{H^1}(h_2, 1)$. So since $\tilde{H}$ contains $N_{H^1}(h_2)$ and $\tilde{H}$ was contracted to h_2 in the construction of H, this yields (β).

Let us now verify that $(H, \mathcal{V})$ satisfies (γ). For this, let $v \in V_{h_2}$. If $v \in B_2$, then $d_G(v, v_2) \leqslant r_2 \leqslant r_2 + 2R'_0(K) + 1$ as desired. Otherwise, by construction, v is contained in V^1_h for some node h of $\tilde{H} \subseteq H^1$. If $h \in B_{H^1}(h_2, 1)$, then $d_G(v, v_2) \leqslant r'_2 + 2R_0(K) + 5K + 1$ by (c). Otherwise, by (b), we have $d_G(v_2, v) \leqslant r_2 + 1 + 2\operatorname{rad}_G(V^1_h) \leqslant r_2 + 1 + 2R'_0(K)$ because V^1_h meets $B_2^{+1} = B_G(B_2, 1)$ as $h \in V(\tilde{H})$. Thus, $(H, \mathcal{V})$ satisfies (γ). Further, (δ) holds for $(H, \mathcal{V})$ by (d).

We are thus left to check (ε). For this, let D' be a component of $G - Y$ that meets C and D; in particular, since $N_G(D) \cap B_1 \subseteq V(Y^1)$, we have that D' is a component of

$G-(V(Y^1)\cup B_2)$. Let $\hat{D}$ be the component of $D-Y^1$ that contains D'. Since $D'\cap C\neq\emptyset$, there is a $\hat{D}$–B_1 path Q in C; let q be its first vertex in Y^1; in particular, $q\in N_G(\hat{D})$. Then

$$d_G(q,v_2)\geqslant d_G(C,v_2)>r_2\geqslant r_2'+\ell(K)=r_2'+2\cdot R_0(K)+5K+2,$$

which by (c) implies that $q\notin V_h^1$ for every $h\in B_{H^1}(h_2,1)$. By (e), there is a bag V_h^1 of $(H^1,\mathcal{V}^1)$ such that $q\in N_G(\hat{D})\subseteq V_h^1$; by the previous observation, we have $h\notin B_{H^1}(h_2,1)$, and hence $\operatorname{rad}_G(V_h^1)\leqslant R_0'(K)$ as we have shown in the proof of (β).

Now observe that

$$N_G(D')=(N_G(D')\cap V(Y^1))\cup(N_G(D')\cap B_2)\subseteq N_G(\hat{D})\cup(N_G(D')\cap B_2)\tag{14.6}$$

where we used that $D'\subseteq\hat{D}-B_2$. Since $D'\subseteq\hat{D}$, $N_G(\hat{D})\subseteq V_h^1$ and $v_2\notin V(\hat{D})$, every D'–v_2 path meets V_h^1. 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_h^1. Since also $N_G(\hat{D})\subseteq V_h^1$, we have by (14.6) for all $u,w\in N_G(D')$ that

$$d_G(u,w)\leqslant d_G(u,V_h^1)+2\cdot\operatorname{rad}_G(V_h^1)+d_G(V_h^1,w)\leqslant\ell(K)+2\cdot R_0'(K)+\ell(K)\leqslant R_2(K)-1.$$

This shows that $(H,\mathcal{V})$ satisfies (ε) and thus concludes the proof. □

Proof of Lemma 14.5.2. Suppose that G has no K-fat K_4 minor for $K\in\mathbb{N}_{\geqslant1}$; otherwise, we are done. Since $4\cdot(2(\ell(K)+22K+1)+11K+2)=R_1(K)$, by Corollary 14.5.10 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 $d_G(G[D,1],P)\geqslant5K$. Let $\mathcal{D}$ be the set of components D of $G-(B_1\cup B_2')$ which meet C and attach to B_2'. We note that these also attach to B_1, as $w'\in V(C^*)\subseteq V(C)$ and C is connected. Set $B_2:=B$.

Note that $r>\ell(K)$, since there are two vertices in $\partial_G C^*$ which are at least $R_1(K)>2(\ell(K)+22K+1)$ apart. We may thus apply Lemma 14.5.11 to $v_1:=w'$, $v_2:=w$, $r_1:=22K$, $r_2:=r$ and every component $D\in\mathcal{D}$. This yields for every such D an honest partial graph-decomposition $\mathcal{H}^D=(H^D,\mathcal{V}^D)$ of G with support $Y^D\subseteq G[B_1\cup V(D)\cup B_2]$ modelled on a graph $H^D\in\mathcal{H}_{SP}$ with terminals h_1^D,h_2^D which satisfies (α) to (ε).

We define the desired decomposition $(H,\mathcal{V})$ in multiple steps and check its desired properties afterwards. First, we define a decomposition $(H^1,\mathcal{V}^1)$ of $Y^1:=G[\bigcup_{D\in\mathcal{D}}V(Y^D)]\subseteq G[B_1\cup(\bigcup_{D\in\mathcal{D}}V(D))\cup B_2]\subseteq G[V(C)\cup B_2]$. The graph H^1 is obtained from the disjoint union of the H^D and three new vertices h_1,h_2,g by identifying all h_1^D with h_1 as well as all h_2^D with h_2 and adding the edge h_2g. Then we assign the nodes of H^1 bags $V_g^1:=N_G(C)$, $V_{h_2}^1:=\bigcup_{D\in\mathcal{D}}V_{h_2^D}^D\supseteq B_2$, and $V_{h_1}^1:=\bigcup_{D\in\mathcal{D}}V_{h_1^D}^D=B_1$ and $V_h^1:=V_h^D$ for all other $h\in V(H^1)$ where D is the unique component in $\mathcal{D}$ with $h\in V(H^D)$. Then $(H^1,\mathcal{V}^1)$ is a graph-decomposition of Y^1. Indeed, it follows from $G[B_1\cup B_2]\subseteq Y^D\subseteq G[B_1\cup V(D)\cup B_2]$ for $D\in\mathcal{D}$ that the Y^D intersect only in B_1 and B_2. Hence, $(H^1,\mathcal{V}^1)$ is a graph-decomposition of Y^1, as every $(H^D,\mathcal{V}^D)$ is a graph-decomposition of Y^D and $B_1=V_{h_1^D}^D$ and $B_2\subseteq V_{h_2^D}^D$.

Next, we adjust $(H^1,\mathcal{V}^1)$ to a graph-decomposition $(H,\mathcal{V}^2)$ of $Y^2\supseteq Y^1$: Let C' be a component of $G-Y^1$ which meets C. In particular, C' is contained in $G-(B_1\cup B_2)$ because $B_1,B_2\subseteq V(Y^1)$. Let $D_{C'}$ be the unique component of $G-(B_1\cup B_2')$ which contains C'. By $\mathcal{C}$ we denote the set of all components C' of $G-Y^1$ which meet C and

whose $D_{C'}$ attach to B'_2. In particular, for every $C' \in \mathcal{C}$ we have that $D_{C'} \in \mathcal{D}$ and C' is a component of $G - Y^{D_{C'}}$ that meets C and $D_{C'}$, since $Y^{D_{C'}} \subseteq Y^1$. Hence, it follows from (ε) of $(H^{D_{C'}}, \mathcal{V}^{D_{C'}})$ that there is a node $h_{C'} \in V(H^{D_{C'}})$ such that $N_G(C') \subseteq V_{h_{C'}}^{D_{C'}}$ and $\mathrm{rad}_G(N_G(C')) \leqslant R_2(K) - 1$.

Now we obtain H from H^1 by adding for each component $C' \in \mathcal{C}$ a new node $h'_{C'}$ and the edge $h_{C'}h'_{C'}$. We assign an $h'_{C'}$ the bag $V^2_{h'_{C'}} := N_G(C') \cup \partial_G C'$ for $C' \in \mathcal{C}$. Moreover, we set $V^2_h := V^1_h$ for all old nodes h of $H^1 \subseteq H$. Then it is immediate from the construction that $(H, \mathcal{V}^2)$ is a graph-decomposition of $Y^2 := G[V(Y^1) \cup (\bigcup_{C' \in \mathcal{C}} \partial_G(C'))]$. We now obtain our final graph-decomposition $(H, \mathcal{V})$ of $Y := Y^2 \cap G[C, 1]$ as the restriction of $(H, \mathcal{V}^2)$ to Y.

It remains to check that $(H, \mathcal{V})$ has the desired properties. By definition, $Y \subseteq G[C, 1]$. We also have $\partial_G C \subseteq V(Y)$. Indeed, every vertex $v \in \partial_G C \subseteq N_G(B_2)$ which was not already contained in $V(Y^1) \supseteq B_2$ lies in $\partial_G C'$ for some component C' of $G - Y^1$. Then $C' \in \mathcal{C}$, i.e. the component $D_{C'}$ of $G - (B_1 \cup B'_2)$ attaches to B'_2: because $v \in V(C') \subseteq V(D_{C'})$ and B_2 and B_1 are disjoint, every shortest v–$N_G(B'_2)$ path in G is through $B_2 \setminus B'_2$ and hence contained in $D_{C'}$. Now by construction, $v \in V_{h'_{C'}}$. Moreover, the graph H is in $\mathrm{Forb}_{\preccurlyeq}(K_4)$: The graph H^1 is obtained from the $H^D \in \mathcal{H}_{SP}$ by parallel compositions of the graphs $H^D \in \mathcal{H}_{SP}$ followed by a series composition with a path $h_2 g$ with terminals h_2 and g; so H^1 is again in $\mathcal{H}_{SP}$ with terminals h_1 and g by Proposition 14.5.1. 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 $V_g = V^2_g = V^1_g = N_G(C)$ 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 $V^1_h = V^2_h$, and hence $\mathrm{rad}_G(V^2_h) \leqslant \mathrm{rad}_G(V^1_h)$. By construction of $(H^2, \mathcal{V}^2)$, every node in $h \in H_2 - H_1$ is of the form $h'_{C'}$ for some $C' \in \mathcal{C}'$, so their bags $V^2_{h'_{C'}}$ are contained in the respective $V^1_{h_{C'}} \cup N_G(V^1_{h_{C'}})$, and hence $\mathrm{rad}_G(V^2_{h'_{C'}}) \leqslant \mathrm{rad}_G(V^1_{h_{C'}}) + 1$. Thus, as $\mathrm{rad}_G(V_{h'_{C'}}) = \mathrm{rad}_G(V^2_{h'_{C'}}) \leqslant R_2(K)$, the graph-decomposition $(H, \mathcal{V})$ satisfies (ii). Also $(H, \mathcal{V})$ satisfies (iii): The $(H^D, \mathcal{V}^D)$ have radial spread at most 3. Thus, $(H^1, \mathcal{V}^1)$ has radial spread at most 6 by construction. This yields that $(H, \mathcal{V}^2)$ has radial spread at most 7 which yields that its restriction $(H, \mathcal{V})$ also has radial spread at most 7.

We claim that $(H, \mathcal{V})$ satisfies (iv). Indeed, let $\tilde{C}$ be a component of $G - Y$. First, assume that $\tilde{C}$ does not meet C. Since $N_G(C) \subseteq V(Y) \subseteq B_G(C, 1)$, its neighbourhood $N_G(\tilde{C})$ is contained in $N_G(C) = V_g$. Also, we have $\mathrm{rad}_G(N_G(C)) \leqslant \mathrm{rad}_G(B) \leqslant r$, since C is a component of $G - B$ and $B = B_G(w, r)$. So we may assume from now on that $\tilde{C}$ meets C. Since $N_G(C), B_1 \subseteq V(Y)$, the component $\tilde{C}$ is contained in $G - (B_1 \cup B_2)$, and thus in C. As $Y = Y^2 \cap G[C, 1]$, $Y^2 \supseteq Y^1$ and $\tilde{C} \subseteq C$, there is a unique component C' of $G - Y^1$ which contains $\tilde{C}$. If $D_{C'}$ does not attach to B'_2, its neighbourhood $N_G(D_{C'})$ is contained in $B_1 = V^1_{h_1} = V^2_{h_1} = V_{h_1}$, and thus $D_{C'} = C' = \tilde{C}$ by their respective definition. So we may assume that $D_{C'}$ attaches to B'_2, i.e. $C' \in \mathcal{C}$. Then the construction of $(H, \mathcal{V}^2)$ ensured that $V_{h'_{C'}} = V^2_{h'_{C'}} = N_G(C') \cup \partial_G C'$. Also $\tilde{C}$ is a component of $C' - \partial_G C'$, since $Y^2 = G[V(Y^1) \cup \bigcup_{C' \in \mathcal{C}} \partial_G C']$. Thus, we have $N_G(\tilde{C}) \subseteq \partial_G C' \subseteq V^2_{h'_{C'}} = V_{h'_{C'}}$. We have already seen above that $\mathrm{rad}_G(V_{h'_{C'}}) \leqslant R_2(K)$. □

CHAPTER 15

A Menger-type theorem for two induced paths

This chapter is based on [AHJKW24] and joint work with Sandra Albrechtsen, Tony Huynh, Paul Knappe and Paul Wollan.

We give an approximate Menger-type theorem for when a graph G contains two $X-Y$ paths P_1 and P_2 such that $P_1 \cup P_2$ is an induced subgraph of G. More generally, we prove that there exists a function $f(d) \in O(d)$, such that for every graph G and $X, Y \subseteq V(G)$, either there exist two $X-Y$ paths P_1 and P_2 such that $d_G(P_1, P_2) \geqslant d$, or there exists $v \in V(G)$ such that the ball of radius $f(d)$ centered at v intersects every $X-Y$ path.

15.1 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 $d_G(u,v)$, is the length of a shortest $u-v$ path in G. The *ball of radius r around v* is $B_G(v,r) := \{u \in V(G) : d_G(u,v) \leqslant r\}$. 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 $d_G(H_1, H_2)$, 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 [M27]. 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 [B92; B91]. 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 15.1.1. *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 15.1.1 immediately follows from a more general result where we ask that the two paths be at distance at least d.

Theorem 40. *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 $d_G(P_1, P_2) \geqslant d$ 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 40 (with a slightly worse constant of 272) [GP25].

We conjecture that Theorem 40 should generalize to an arbitrary number of $X - Y$ paths at pairwise distance d.

Conjecture 41. *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 $d_G(P_i, P_j) \geqslant d$ 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 41 actually implies the general case.

Theorem 15.1.2 (McCarty and Seymour, 2023)**.** *If Conjecture 41 holds for $d = 3$ with constant c, then it holds for all $d \geqslant 3$ with constant $3c$.*

Proof. Let H be the dth power of G, that is, $V(H) = V(G)$, and u and v are adjacent in H if and only if $d_G(u, v) \leqslant d$. If there exist $k - 1$ balls of radius $3c$ in H whose union intersects all $X - Y$ paths in H, then there are $k - 1$ balls of radius $3cd$ in G so that their union intersects all $X - Y$ paths in G. Otherwise (by assumption), there are $X - Y$ paths $P_1, \ldots, P_k$ in H such that $d_H(P_i, P_j) \geqslant 3$ for all distinct i, j. Suppose $P_1 = v_0 \ldots v_t$. For each $0 \leqslant i \leqslant t - 1$, there is a path S_i of length at most d in G between v_i and v_{i+1}. Thus, there is a $v_0 - v_t$ path P_1^* in G such that $P_1^* \subseteq \bigcup_{i=0}^{t-1} S_i$. We define $P_2^*, \ldots P_k^*$ similarly. We claim that $d_G(P_i^*, P_j^*) \geqslant d$ for all distinct i, j. Towards a contradiction, suppose $d_G(P_i^*, P_j^*) \leqslant d$ for some $i < j$. Then there is a $V(P_i) - V(P_j)$ path in G of length at most $\frac{d}{2} + d + \frac{d}{2} \leqslant 2d$. Hence, $d_H(P_i, P_j) \leqslant 2$, which is a contradiction. □

We would like to point out that Nguyen, Scott, and Seymour [NSS25a] have recently found a counterexample to Conjecture 41.

Our proof shows that we may take $c = 129$ in Theorem 40. On the other hand, the following lemma shows that c must be at least $3/2$.

Lemma 15.1.3. *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 $d_G(P_1, P_2) \geqslant d$, 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 15.1. If we let $X = \{x_1, x_2, x_3\}$ and $Y = \{y_1, y_2, y_3\}$, then there do not exist two disjoint $X - Y$ paths P_1 and P_2 such that $d_H(P_1, P_2) \geqslant 2$.

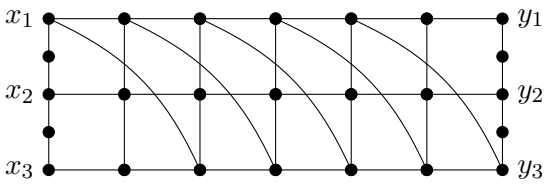

Figure 15.1: A graph without two disjoint anti-complete $\{x_1, x_2, x_3\} - \{y_1, y_2, y_3\}$ paths and no ball of radius one hitting all $X - Y$ paths.

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 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 $P = v_0v_1 \dots v_n$ and $0 \leqslant i \leqslant j \leqslant n$, we write $v_iP := v_iv_{i+1} \dots v_n$, $Pv_i := v_0v_1 \dots v_i$ and $v_iPv_j := v_iv_{i+1} \dots v_j$ for the corresponding subpaths of P.

15.2 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 15.2 for a visualisation.

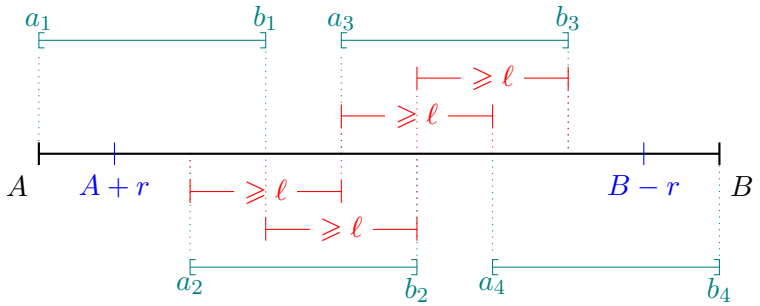

Figure 15.2: A clean (A, B, r)-system $([a_i, b_i])_{i\in[4]}$ indicated in teal which is interlaced with buffer ℓ.

Definition 15.2.1. 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 15.3).

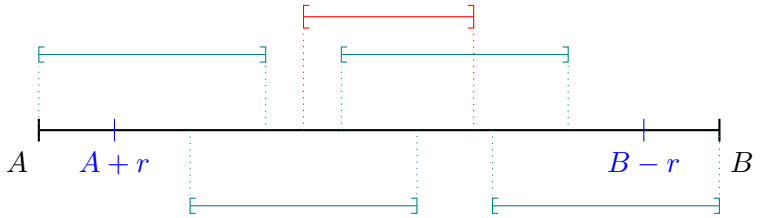

Figure 15.3: The (A, B, r)-system indicated in teal is clean. However, adding the red interval gives an (A, B, r)-system which is not clean.

Lemma 15.2.2. *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 15.2.3. *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. □

15.3 Proof of Theorem 40

We now prove our main theorem, which we restate for convenience.

Theorem 40. *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 $d_G(P_1, P_2) \geqslant d$ or there exists $z \in V(G)$ such that $B_G(z, cd)$ intersects every $X - Y$ path.*

Proof. Let $d_1 = 5d$ and $d_2 = 62d$. If G has no $X - Y$ path, we may take z to be any vertex of G. Thus, let $P = v_0 \dots v_n$ be a shortest $X - Y$ path in G. Let $\mathcal{H}$ be the family of components of $G - B_G(P, d_1)$. If some $H \in \mathcal{H}$ contains an $X - Y$ path W, then we may take $P_1 = P$ and $P_2 = W$. Note that $d_G(P_1, P_2) \geqslant d_1 \geqslant d$. Thus, we may assume that no $H \in \mathcal{H}$ contains vertices of both X and Y.

For each $H \in \mathcal{H}$, we define $a(H)$ to be $-d_2$ if H contains a vertex of X. Otherwise, $a(H)$ is the smallest index j such that $d_G(v_j, H) = d_1 + 1$. Similarly, we define $b(H)$ to be $n + d_2$ if H contains a vertex of Y. Otherwise, $b(H)$ is the largest index j such that $d_G(v_j, H) = d_1 + 1$. For each $H \in \mathcal{H}$, we define $I(H)$ to be the interval $[a(H), b(H)]$ (see Figure 15.4). 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 15.2.3, 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 $Z = [j - 2d_2, j + 2d_2] \cap (-d_2, d_2 + n)$.

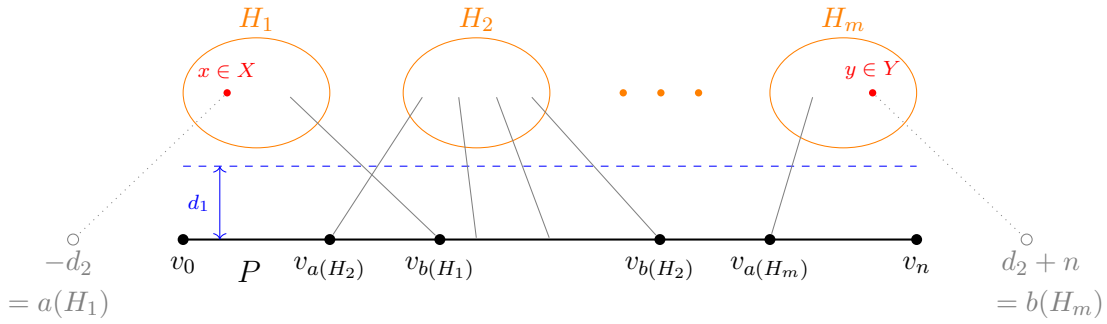

Figure 15.4: The components H_i along with their respective $a(H_i)$ and $b(H_i)$ represented by the corresponding vertex in P or a vertex in X or Y.

Suppose the latter holds. We claim that $B_G(z, cd)$ intersects all the $X - Y$ paths for $z := v_j$. Suppose not and let Q be an $X - Y$ path which avoids $B_G(z, cd)$. Let $R_1 := Pv_{j-2d_2-1}$ and $R_2 := v_{j+2d_2+1}P$. Since $cd \geqslant 2d_2 + d_1$, the path Q in particular avoids $B_G(v_{j-2d_2}Pv_{j+2d_2}, d_1) \supseteq B_G(P, d_1) \setminus (B_G(R_1, d_1) \cup B_G(R_2, d_1))$.

For $i \in \{1, 2\}$ let $\mathcal{H}_i$ be those components $H \in \mathcal{H}$ which have a neighbour in $B_G(R_i, d_1)$. By assumption, no interval in $I(\mathcal{H}')$ intersects both segments of $[-d_2, d_2 + n] \setminus Z$. Since $\mathcal{H}' \subseteq \mathcal{H}$ are the components which correspond to the maximal intervals in $I(\mathcal{H})$, we conclude that no interval in $I(\mathcal{H})$ intersects both segments of $[-d_2, d_2 + n] \setminus Z$. For each $H \in \mathcal{H}$, the choice of $a(H)$ and $b(H)$ implies that if v_h is a vertex of P such that $d_G(v_h, H) = d_1 + 1$, then $h \in I(H) = [a(H), b(H)]$. Therefore, we conclude that $\mathcal{H}_1$ and $\mathcal{H}_2$ are disjoint, no $H \in \mathcal{H}_1$ meets Y, and no $H \in \mathcal{H}_2$ meets X.

Since P is a shortest path, $d_G(v_k, v_\ell) = |k - \ell|$ for all $0 \leqslant k, \ell \leqslant n$. Hence, $d_G(R_1, R_2) = 4d_2 + 2$. Since an edge in G from $B_G(R_1, d_1)$ to $B_G(R_2, d_1)$ would witness that $d_G(R_1, R_2) \leqslant 2d_1 + 1$ but $2d_1 + 1 < 4d_2 + 2$, it follows that no such edge exists.

Altogether, the $X - Y$ path Q avoiding $B_G(z, cd)$ must be contained in $B_G(R_i, d_1) \cup \bigcup \mathcal{H}_i$ for some $i \in \{1, 2\}$. We consider the case $i = 1$ as the case $i = 2$ is analogous. Since no $H \in \mathcal{H}_1$ meets Y, the 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 $H_1, \ldots, H_m$ of a subset of $\mathcal{H}'$ such that $I(H_1), \ldots, I(H_m)$ is an interlaced system with buffer d_2. By Lemma 15.2.2, we may also assume that $I(H_1), \ldots, I(H_m)$ 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 $C_i \notin \{H_1, \ldots, H_m\}$ 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 $\mathcal{M} := \{H_1, \ldots, H_m\}$ and $\mathcal{S} := \{C_1, \ldots, C_s\}$ 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 $H_1, \ldots, H_m$ and of $C_1, \ldots C_s$. Rename

the corresponding sequence $H'_1, \dots H'_{m+s}$.

Claim 1. *$d_G(v_{b(H'_j)}, v_{a(H'_{j+1})}) \geqslant 30d$ for every $j \in [m+s-1]$.*

Proof of the claim. If $H'_j, H'_{j+1} \in \mathcal{M}$, then the claim follows directly from the choice of $\mathcal{S}$. Thus, we may assume that $H'_j \in \mathcal{M}$ and $H'_{j+1} \in \mathcal{S}$, since the case $H'_j \in \mathcal{S}$ and $H'_{j+1} \in \mathcal{M}$ is analogous. Then we have that $a(H'_{j+1}) + d_2 \leqslant a(H'_{j+2}) \leqslant b(H'_j)$. Thus, the claim follows from $d_2 \geqslant 30d$. ■

Claim 2. *$d_G(v_{a(H'_j)}, v_{a(H'_{j+2})}) \geqslant d_2$ for every $j \in [m+s-2]$.*

Proof of the claim. If at least two of $\{H'_j, H'_{j+1}, H'_{j+2}\}$ are in $\mathcal{M}$, then the claim follows from the choice of $\mathcal{M}$. Otherwise, $H'_j \in \mathcal{S}$ and $H'_{j+2} \in \mathcal{S}$. Then the claim follows from the choice of $\mathcal{S}$. ■

Claim 3. *If $q \geqslant p+5$, then $d_G(H'_p, H'_q) \geqslant d_2 - 2d_1 + 2 - 30d$. Moreover, if $q \geqslant p+7$, then $d_G(H'_p, H'_q) \geqslant d_2 - 2d_1 + 2$.*

Proof of the claim. 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

$$d_G(v_\ell, v_r) \leqslant d_G(v_\ell, h'_p) + d_G(h'_p, h'_q) + d_G(h'_q, v_r) = 2d_1 + (||Q|| - 2).$$

Suppose there are at least four elements of $\mathcal{M}$ among $H'_p, \dots, H'_q$. Since $(I(H_i))_{i=1}^k$ is a clean interlaced system with buffer d_2, it follows that $r - \ell \geqslant d_2$. Since P is a shortest $X-Y$ path, we thus obtain $d_G(v_\ell, v_r) = r - \ell \geqslant d_2$. Together with $d_G(v_\ell, v_r) \leqslant 2d_1 + (||Q|| - 2)$, this yields $||Q|| \geqslant d_2 - 2d_1 + 2$.

Note that if $q \geqslant p+7$, then there are at least four elements of $\mathcal{M}$ among $H'_p, \dots, H'_q$. Thus, we may assume that $q \in \{p+5, p+6\}$ and that are at most three elements of $\mathcal{M}$ among $H'_p, \dots, H'_q$. Note that this can only happen if the sequence $H'_p, \dots, H'_q$ alternates between elements in $\mathcal{M}$ and $\mathcal{S}$. By symmetry, we may assume that $H'_p \in \mathcal{S}$. We have $\ell \leqslant b(H'_p) \leqslant b(H'_{p+1})$ and $a(H'_{p+5}) \leqslant a(H'_q) \leqslant r$. Since $H'_{p+4} \in \mathcal{S}$, $a(H'_{p+4}) \leqslant a(H'_{p+5}) - d_2$. Since $H'_{p+2} \in \mathcal{S}$, it follows that $d_G(v_{b(H'_{p+1})}, v_{a(H'_{p+3})}) \leqslant 30d$. Therefore, $d_G(v_\ell, v_r) \geqslant d_2 - 30d$. On the other hand, $d_G(v_\ell, v_r) \leqslant 2d_1 + (||Q|| - 2)$, as above. Thus, we have $||Q|| \geqslant d_2 - 2d_1 + 2 - 30d$, as required. ■

Let $H' = \bigcup_{i=1}^{m+s} H'_i$. A *fruit tree* is a subgraph $\mathcal{W} := \bigcup_{i \in I} H'_i \cup W_1 \cup \dots \cup W_k$ of G (see Figure 15.5) 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 $v \in V(C)$, $d_{\mathcal{W}}(v, V(H')) \leqslant \ell d$.*

Proof of the claim. 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, $d_{C_1}(w_1, V(H')) \leqslant \ell_1 d$ and $d_{C_2}(w_2, V(H')) \leqslant \ell_2 d$. Therefore, we are done unless $d_{W_{i_\ell}}(v, w_1) > (\ell - \ell_1)d$ and $d_{W_{i_\ell}}(v, w_2) > (\ell - \ell_2)d$. But this implies $||W_{i_\ell}|| > (\ell - \ell_1)d + (\ell - \ell_2)d \geqslant (\ell + 1)d \geqslant (\max\{\ell_1, \ell_2\} + 2)d$, which contradicts that $\mathcal{W}$ is d-small. ■

Claim 5. *Let $\mathcal{W}$ be a d-small fruit tree with composite paths $W_1, \dots, W_k$. Then each component of $\bigcup_{i=1}^{k} W_i$ consists of at most 4 composite paths.*

Proof of the claim. 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

$$d_C(u_p, u_q) \leqslant \sum_{i=1}^{c_1}(i+1)d + (4+2)d + \sum_{i=1}^{c_2}(i+1)d$$

by joining u_p to $W_{i_{\ell'}}$ in C_1, walking along $W_{i_{\ell'}}$ and then joining $W_{i_{\ell'}}$ to u_q in C_2.

If $c \geqslant 7$, then it follows from $c_1, c_2 \leqslant 4$ that

$$d_C(u_p, u_q) \leqslant \sum_{i=1}^{4}(i+1)d + (4+2)d + \sum_{i=1}^{4}(i+1)d = 34d < d_2 - 2d_1 + 2.$$

However, this contradicts Claim 3.

If $c \in \{5, 6\}$, then it follows from $c_1 + c_2 + 1 = c \leqslant 6$ that

$$d_C(u_p, u_q) \leqslant \sum_{i=1}^{4}(i+1)d + (4+2)d + \sum_{i=1}^{1}(i+1)d = 22d < d_2 - 2d_1 + 2 - 30d.$$

Again, this is a contradiction to Claim 3. ■

For each fruit tree $\mathcal{W} = \bigcup_{i\in I} H'_i \cup W_1 \cup \cdots \cup W_k$, we 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 15.5.

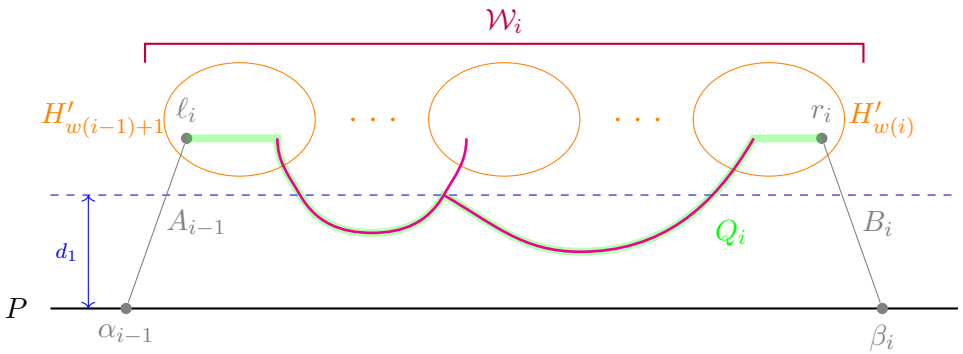

Figure 15.5: The fruit tree $\mathcal{W}_i$ of the orchard with its corresponding components (orange) and its composite paths (magenta).

We will now recursively define two $X - Y$ walks P_1 and P_2 such that $d_G(P_1, P_2) \geqslant d$. These walks then contain the desired paths. In order to show that our construction indeed yields $d_G(P_1, P_2) \geqslant d$, we will need to ensure that $d_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$.

Initialize $P_1 := P\alpha_1$ and $P_2 := \ell_1 Q_1 r_1$. Now suppose that we have already defined P_1 and P_2 such that, for some $i \in [t]$,

(i) P_1 starts in v_0 and ends in α_i, and P_2 starts in ℓ_1 and ends in r_i; or P_1 starts in v_0 and ends in r_i, and P_2 starts in ℓ_1 and ends in α_i,

(ii) $P_1 \cup P_2$ does not contain any vertices of $\bigcup_{j=w(i)+1}^{m+s} H'_j$,

(iii) among all components of $(P_1 \cup P_2) \cap P$, the component P' containing α_i occurs last on P (note that α_i could be the left or right 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 $d_P(\alpha_i, \beta_j) \leqslant 20d$ and if additionally $d_P(\alpha_i, \beta_j) > 10d$, then α_{i+1} occurs between α_i and β_j along P, and

(vi) $d_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$.

We maintain these properties throughout the construction. Note that property (vi) suffices in order to prove that the final P_1 and P_2 have distance at least d in G. The other properties ensure that we can update the paths in the induction step without losing property (vi).

Now we describe how to extend P_1 and P_2. Suppose P_1 ends in α_i, and P_2 ends in r_i (the construction is analogous in the other case). Further, let β_j be the other 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 $d_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$ 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 15.6).

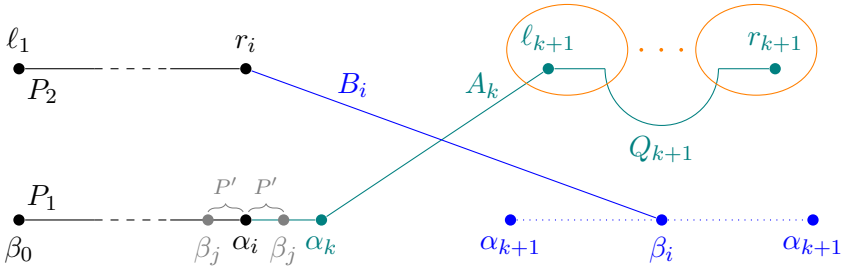

Figure 15.6: The old paths P_1 and P_2 in black together with their updates in teal and blue, respectively. Note that α_{k+1} may occur before or after β_i and that β_j may occur before or after α_i.

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 $d_P(\beta_i, \{\alpha_i, \beta_j\}) \geqslant 30d - 20d = 10d$ by assumption and Claim 1, it suffices to ensure that $d_P(\beta_i, \alpha_k) \geqslant 10d$ if $k \neq i$ and that $d_P(\alpha_{k+1}, \{\alpha_k, \beta_j\}) \geqslant 10d$ in order to obtain (vi).

To find the suitable $k \in \{i, i+1, i+2, i+3\}$ for the update of P_1 and P_2, we consider where $\alpha_i, \alpha_{i+1}, \alpha_{i+2}, \alpha_{i+3}$ occur relative to β_i on P. We depict the cases in Figures 15.7 to 15.10. In each of the figures, we indicate the corresponding assumptions in orange and their implications in green.

Suppose α_{i+1} occurs after β_i along P. Then $d_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) \geqslant d_P(\beta_i, \{\alpha_i, \beta_j\}) \geqslant 10d$, and we set $k := i$ (see Figure 15.7). By assumption, (v) and (vi) are clear.

Now suppose that α_{i+1} occurs before β_i along P. Further, suppose α_{i+2} occurs after β_i along P (see Figure 15.8). In this case, we proceed as follows. If $d_P(\alpha_{i+1}, \beta_i) \geqslant 10d$, then we set $k := i+1$ (see Figure 15.8a). Again, (v) and (vi) are clear from the assumption. We

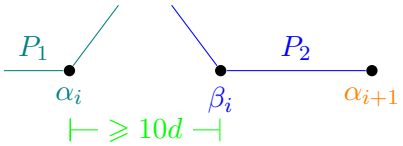

Figure 15.7: α_{i+1} occurs after β_i along P.

remark that by assumption this in particular is the case if $d_P(\alpha_i, \beta_j) > 10d$ and β_j occurs after α_i along P, since then α_{i+1} occurs between α_i and β_j along P and, as shown above, $d_P(\beta_i, \beta_j) \geqslant 10d$.

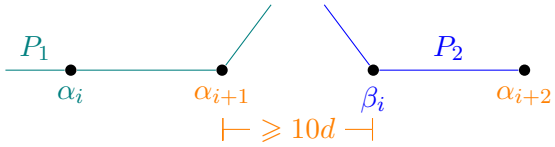

(a) $d_P(\alpha_{i+1}, \beta_i) \geqslant 10d$

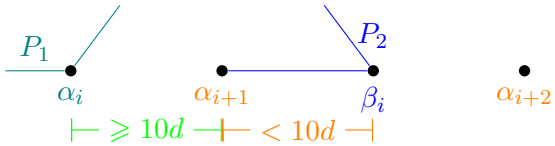

(b) $d_P(\alpha_{i+1}, \beta_i) < 10d$

Figure 15.8: α_{i+1} occurs before β_i along P and α_{i+2} occurs after β_i along P.

Otherwise, $d_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) \geqslant 10d$. Indeed, either β_j occurs before α_i along P and thus we have $d_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) = d_P(\alpha_{i+1}, \alpha_i) = d_P(\alpha_i, \beta_i) - d_P(\alpha_{i+1}, \beta_i) \geqslant 30d - 10d \geqslant 10d$ by Claim 1, or β_j occurs after α_i along P and thus we have $d_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) = d_P(\alpha_{i+1}, \beta_j) = d_P(\alpha_i, \beta_i) - d_P(\alpha_{i+1}, \beta_i) - d_P(\alpha_i, \beta_j) \geqslant 30d - 10d - 10d \geqslant 10d$. In this case we set $k := i$ (see Figure 15.8b). Then (vi) holds by the previous calculation. Moreover, recall that $d_P(\alpha_{k+1}, \beta_i) < 10d$ by assumption, and thus also (v) holds.

We may thus assume that both α_{i+1} and α_{i+2} occur before β_i along P. Suppose further that α_{i+3} occurs after β_i along P (see Figure 15.9). In this case, we proceed as follows. If $d_P(\alpha_{i+2}, \beta_i) \geqslant 10d$, then we set $k := i + 2$ (see Figure 15.9a). Then (v) and (vi) are clear from the assumption.

Otherwise, if $d_P(\alpha_{i+2}, \beta_i) < 10d$, we have $d_P(\alpha_{i+2}, \{\alpha_i, \beta_j\}) \geqslant d_P(\alpha_{i+2}, \alpha_i) - 20d \geqslant d_2 - 20d \geqslant 10d$ by assumption and Claim 2. If $d_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant 10d$, then we set $k := i + 1$ (see Figure 15.9b). Then (vi) holds by assumption. Since also $d_P(\alpha_{k+1}, \beta_i) < 10d$ by assumption, (v) holds too.

Otherwise, $d_P(\alpha_{i+1}, \beta_i) = d_P(\alpha_{i+1}, \alpha_{i+2}) + d_P(\alpha_{i+2}, \beta_i) \leqslant 20d$ and $d_P(\alpha_{i+1}, \{\alpha_i, \beta_j\}) \geqslant d_P(\alpha_i, \alpha_{i+2}) - d_P(\alpha_{i+1}, \alpha_{i+2}) - 20d \geqslant d_2 - 10d - 20d \geqslant 10d$ by assumption and Claim 2, and we then set $k := i$ (see Figure 15.9c). 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 $d_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant d_2$. Indeed, the fact that $\alpha_{i+3} =$

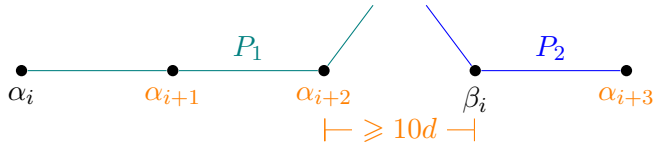

(a) $d_P(\alpha_{i+2}, \beta_i) \geqslant 10d$

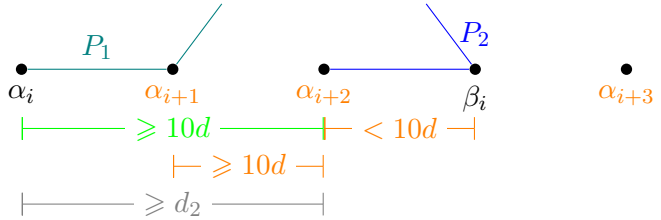

(b) $d_P(\alpha_{i+2}, \beta_i) < 10d$ and $d_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant 10d$

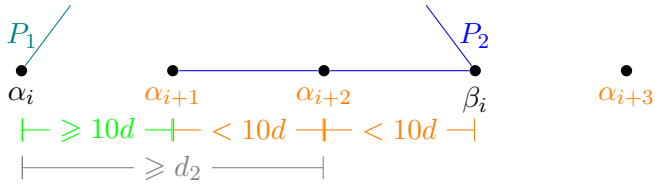

(c) $d_P(\alpha_{i+2}, \beta_i) < 10d$ and $d_P(\alpha_{i+1}, \alpha_{i+2}) < 10d$

Figure 15.9: α_{i+1} and α_{i+2} occur before β_i along P and α_{i+3} occurs after β_i along P.

$v_{a(H'_{w(i+3)+1})}$ occurs before $\beta_i = v_{b(H'_{w(i)})}$ implies that no three of $H'_{w(i)}, \dots, H'_{w(i+3)+1}$ are in $\mathcal{M}$ since $\mathcal{M}$ is clean. Note that $w(i+3)+1 \geqslant w(i)+4$, since the $w(i)$ are strictly increasing by definition. By construction of $(H'_i)_{i\in[m+s]}$, we have that if $H'_i \in \mathcal{S}$ then $H'_{i+1} \in \mathcal{M}$. Hence, $w(i+3)+1 = w(i)+4$, and $H'_{w(i)}, H'_{w(i)+2}, H'_{w(i)+4} \in \mathcal{S}$, and $H'_{w(i)+1}, H'_{w(i)+3} \in \mathcal{M}$. In particular, $\mathcal{W}^{i+1} = H'_{w(i)+1}, \mathcal{W}^{i+2} = H'_{w(i)+2}$ and $\mathcal{W}^{i+3} = H'_{w(i)+3}$ consist of single components of H'. It follows that $d_P(\alpha_{i+1}, \alpha_{i+2}) = d_P(v_{a(H'_{w(i)+2})}, v_{a(H'_{w(i)+3})}) \geqslant d_2$ by the choice of $\mathcal{S}$. Moreover, as $H'_{w(i)+4} \in \mathcal{S}$, we have $H'_{w(i)+5} \in \mathcal{M}$. Together with $H'_{w(i)+1}, H'_{w(i)+3} \in \mathcal{M}$, this implies that $a(H'_{w(i+4)+1}) \geqslant a(H'_{w(i)+5}) \geqslant b(H'_{w(i)+1}) \geqslant b(H'_{w(i)})$ as $\mathcal{M}$ is clean. Thus, $\alpha_{i+4} = v_{a(H_{w(i+4)+1})}$ occurs after $\beta_i = v_{b(H'_{w(i)})}$ along P.

If $d_P(\alpha_{i+2}, \beta_i) \leqslant 20d$, then we set $k := i+1$ (see Figure 15.10a). Since $d_P(\alpha_{i+1}, \alpha_{i+2}) \geqslant d_2 \geqslant 10d$ as shown above, (vi) holds. Moreover, $\alpha_{k+2} = v_{a(H'_{w(k+2)+1})}$ occurs between $\alpha_{k+1} = v_{a(H'_{w(k+1)+1})}$ and β_i by assumption and because $a(H'_{w(k+1)+1}) < a(H'_{w(k+2)+1})$. Thus (v) holds.

If $d_P(\alpha_{i+3}, \beta_i) \geqslant 10d$, then we set $k := i+3$ (see Figure 15.10b). Then (v) and (vi) are clear by assumption.

Otherwise, $d_P(\alpha_{i+2}, \alpha_{i+3}) = d_P(\alpha_{i+2}, \beta_i) - d_P(\alpha_{i+3}, \beta_i) \geqslant 20d - 10d = 10d$. In this case we set $k := i+2$ (see Figure 15.10c). Then (vi) holds by the previous calculation. Moreover, (v) holds since $d_P(\alpha_{i+3}, \beta_i) < 10d$ by assumption.

Claim 6. $d_G(P_1 \cap P, P_2 \cap P) \geqslant 10d$.

Proof of the claim. By construction, we have $d_P(P_1 \cap P, P_2 \cap P) \geqslant 10d$. Since P is a shortest $X - Y$ path, the claim immediately follows. ■

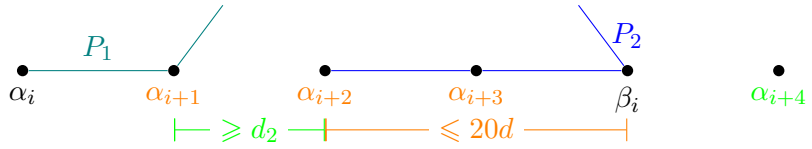

(a) $d_P(\alpha_{i+2}, \beta_i) \leqslant 20d$

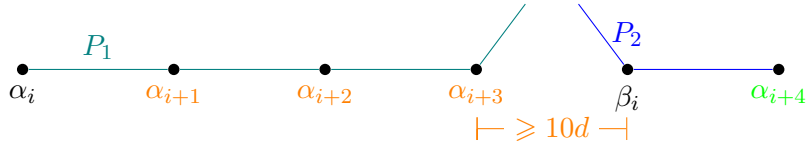

(b) $d_P(\alpha_{i+3}, \beta_i) \geqslant 10d$

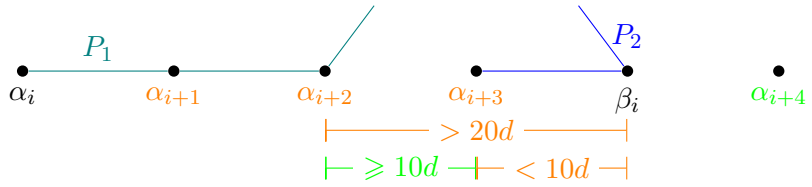

(c) $d_P(\alpha_{i+2}, \beta_i) > 20d$ and $d_P(\alpha_{i+3}, \beta_i) < 10d$

Figure 15.10: α_{i+1}, α_{i+2} and α_{i+3} occur before β_i along P.

Claim 7. $d_G(P_1, P_2) \geqslant d$.

Proof of the claim. Suppose for a contradiction that $p_1 \in V(P_1)$, $p_2 \in V(P_2)$, and $d_G(p_1, p_2) \leqslant d-1$. If $p_1 \in V(Q_i)$ and $p_2 \in V(Q_j)$ for some $i \neq j$, then $d_G(p_1, p_2) \geqslant d$ by the minimality of t. Suppose $p_1 \in V(P)$ and $p_2 \in V(Q_i)$ for some i. By Claim 4 and Claim 5, $d_G(p_2, H') \leqslant 4d$. Since $d_G(P, H') \geqslant d_1$, we have $d_G(p_1, p_2) \geqslant d_1 - 4d \geqslant d$. Suppose $p_1 \in V(P)$ and $p_2 \in V(A_i)$ for some i. By Claim 6, $d_G(p_1, \alpha_i) \geqslant 10d$. This contradicts $d_G(p_1, \alpha_i) \leqslant d_G(p_1, p_2) + d_G(p_2, \alpha_i) \leqslant d - 1 + d_1 < 10d$. The same argument handles the case $p_1 \in V(P)$ and $p_2 \in V(B_i)$ for some i.

Suppose $p_1 \in V(A_i)$ and $p_2 \in V(A_j)$ for some $i < j$. If $d_G(p_1, \ell_i) + d_G(p_2, \ell_j) \leqslant d+1$, then

$$d_G(\ell_i, \ell_j) \leqslant d_G(\ell_i, p_1) + d_G(p_1, p_2) + d_G(p_2, \ell_j) \leqslant (d+1) + (d-1) = 2d,$$

which contradicts the minimality of t. Otherwise,

$$\begin{aligned} d_G(\alpha_i, \alpha_j) &\leqslant d_G(\alpha_i, p_1) + d_G(p_1, p_2) + d_G(p_2, \alpha_j) \\ &= (d_G(\alpha_i, \ell_i) - d_G(\ell_i, p_1)) + d_G(p_1, p_2) + (d_G(\alpha_j, \ell_j) - d_G(\ell_j, p_2)) \\ &\leqslant 2d_1 + 2 - (d+2) + (d-1) < 10d, \end{aligned}$$

which contradicts $d_G(\alpha_i, \alpha_j) \geqslant 10d$ by Claim 6. The same argument handles the case $p_1 \in V(B_i)$ and $p_2 \in V(B_j)$, and $p_1 \in V(A_i)$ and $p_2 \in V(B_j)$, for some $i \neq j$.

Suppose $p_1 \in V(Q_i)$ and $p_2 \in V(A_j)$ for some i, j. Since $d_G(H', P) = d_1 + 1$ and $||A_j|| \leqslant d_1 + 1$,

$$d_G(P, p_2) + d_G(p_2, \ell_j) = ||A_j|| \leqslant d_G(H', P) \leqslant d_G(H', p_1) + d_G(p_1, p_2) + d_G(p_2, P).$$

Moreover, assume that the component of $\mathcal{W}^i \setminus E(H')$ which contains p_1 consists of exactly c composite paths. By Claim 4, $d_{\mathcal{W}^i}(p_1, V(H')) \leqslant cd$. Altogether, we have $d_G(p_2, \ell_j) \leqslant d_G(H', p_1) + d_G(p_1, p_2) \leqslant cd + d - 1$. Therefore,

$$d_G(p_1, \ell_j) \leqslant d_G(p_1, p_2) + d_G(p_2, \ell_j) \leqslant cd + 2d - 2 \leqslant (c+2)d.$$

However, this contradicts the minimality of t, since there is $\mathcal{W}^i - \mathcal{W}^j$ path of length at most $(c+2)d$.

The same argument also handles the case $p_1 \in V(Q_i)$ and $p_2 \in V(B_j)$ for some i, j. By symmetry between p_1 and p_2, there are no more remaining cases. ■

This concludes the proof. □

Acknowledgements

We are grateful to Rose McCarty and Paul Seymour for allowing us to include the proof of Theorem 15.1.2 in our paper [AHJKW24]. We also thank Agelos Georgakopoulos and Panagiotis Papazoglou for informing us of the existence of their preprint [GP25]. We also acknowledge Adrian Thananopavarn for correcting the lowerbound in Lemma 15.1.3 from an earlier version. Finally, we thank the two anonymous referees of our paper [AHJKW24] on which this chapter is based for their valuable input which greatly improved its readability.

Bibliography

[AABKN25] T. Abrishami, S. Albrechtsen, N. Bowler, P. Knappe, and J. K. Nickel. *Locally Interval Graphs Are Circular-Arc Graphs*. 2025. arXiv: 2512.19040.

[AK25] T. Abrishami and P. Knappe. *The Global Structure of Locally Chordal Graphs*. 2025. arXiv: 2512.19044.

[AKK25] T. Abrishami, P. Knappe, and J. Kobler. *Locally Chordal Graphs*. 2025. arXiv: 2501.17320v2.

[AR] T. Abrishami and M. Z. Rácz. "An Algorithmic Perspective on the Local Cover of a Graph". In preparation.

[AL03] F. N. Abu-Khzam and M. A. Langston. "Graph Coloring and the Immersion Order". In: *Computing and Combinatorics*. COCOON 2003. Ed. by T. Warnow and B. Zhu. Red. by G. Goos, J. Hartmanis, and J. Van Leeuwen. Vol. 2697. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 394–403. DOI: 10.1007/3-540-45071-8_40.

[AB09] R. Aharoni and E. Berger. "Menger's Theorem for Infinite Graphs". In: *Inventiones mathematicae* 176.1 (2009), pp. 1–62. DOI: 10.1007/s00222-008-0157-3. arXiv: 0509397.

[A23] S. Albrechtsen. *Optimal Trees of Tangles: Refining the Essential Parts*. 2023. arXiv: 2304.12078.

[A24a] S. Albrechtsen. "Refining Trees of Tangles in Abstract Separation Systems: Inessential Parts". In: *Combinatorial Theory* 4.1 (2024). DOI: 10.5070/C64163854. arXiv: 2302.01808.

[A24b] S. Albrechtsen. *Tangle-Tree Duality in Infinite Graphs*. 2024. arXiv: 2405.06756.

[A25] S. Albrechtsen. "Refining Tree-Decompositions so That They Display the k-Blocks". In: *Journal of Graph Theory* 109.3 (2025), pp. 310–314. DOI: 10.1002/jgt.23230. arXiv: 2403.19585.

[ABJKW24] S. Albrechtsen, H. von Bergen, R. W. Jacobs, P. Knappe, and P. Wollan. *On Vertex Sets Inducing Tangles*. 2024. arXiv: 2411.13656.

[AD25] S. Albrechtsen and J. Davies. *Counterexample to the Conjectured Coarse Grid Theorem*. 2025. arXiv: 2508.15342.

[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*. To appear in *Innovations in Graph Theory*. 2023. arXiv: 2307.08497.

[ADG25] S. Albrechtsen, M. Distel, and A. Georgakopoulos. *Excluding $K_{2,t}$ as a Fat Minor*. 2025. arXiv: 2510.14644.

[ADG26] S. Albrechtsen, M. Distel, and A. Georgakopoulos. *Small Counterexamples to the Fat Minor Conjecture*. 2026. arXiv: 2601.05761.

[AHJKW24] S. Albrechtsen, T. Huynh, R. W. Jacobs, P. Knappe, and P. Wollan. "A Menger-type Theorem for Two Induced Paths". In: *SIAM Journal on Discrete Mathematics* 38.2 (2024), pp. 1438–1450. DOI: 10.1137/23M1573082. arXiv: 2305.04721.

[AJKP24] S. Albrechtsen, R. W. Jacobs, P. Knappe, and M. Pitz. *Linked Tree-Decompositions into Finite Parts*. 2024. arXiv: 2405.06753.

[AJKP25] S. Albrechtsen, R. W. Jacobs, P. Knappe, and M. Pitz. "Counterexamples Regarding Linked and Lean Tree-Decompositions of Infinite Graphs". In: *Journal of Graph Theory* 110.4 (2025), pp. 398–407. ISSN: 1097-0118. DOI: 10.1002/jgt.23279. arXiv: 2405.06755.

[AJKW25] S. Albrechtsen, R. W. Jacobs, P. Knappe, and P. Wollan. "A Characterisation of Graphs Quasi-isometric to K_4-Minor-Free Graphs". In: *Combinatorica* 45.6 (2025), p. 61. DOI: 10.1007/s00493-025-00189-5. arXiv: 2408.15335.

[AFN96] K. Ando, S. Fujishige, and T. Nemoto. "Decomposition of a Bidirected Graph into Strongly Connected Components and Its Signed Poset Structure". In: *Discrete Applied Mathematics* 68.3 (1996), pp. 237–248. DOI: 10.1016/0166-218X(95)00068-3.

[AK04] G. Appa and B. Kotnyek. "Rational and Integral K-Regular Matrices". In: *Discrete Mathematics* 275.1 (2004), pp. 1–15. DOI: 10.1016/S0012-365X(03)00095-5.

[AK06] G. Appa and B. Kotnyek. "A Bidirected Generalization of Network Matrices". In: *Networks* 47.4 (2006), pp. 185–198. DOI: 10.1002/net.20108.

[BG09] J. Bang-Jensen and G. Z. Gutin. *Digraphs*. Springer Monographs in Mathematics. London: Springer, 2009. DOI: 10.1007/978-1-84800-998-1.

[BD08] G. Bell and A. Dranishnikov. "Asymptotic Dimension". In: *Topology and its Applications* 155.12 (2008), pp. 1265–1296. DOI: 10.1016/j.topol.2008.02.011.

[BD02] P. Bellenbaum and R. Diestel. "Two Short Proofs Concerning Tree-Decompositions". In: *Combinatorics, Probability and Computing* 11.6 (2002), pp. 541–547. DOI: 10.1017/S0963548302005369.

[BFGR17] R. Belmonte, F. V. Fomin, P. A. Golovach, and M. S. Ramanujan. "Metric Dimension of Bounded Tree-length Graphs". In: *SIAM Journal on Discrete Mathematics* 31.2 (2017), pp. 1217–1243. DOI: 10.1137/16M1057383. arXiv: 1602.02610.

[BR22] O. Bendele and D. Rautenbach. "Additive Tree $O(\varrho \log n)$-Spanners from Tree Breadth ϱ". In: *Theoretical Computer Science* 914 (2022), pp. 39–46. DOI: 10.1016/j.tcs.2022.02.011.

[BS24] E. Berger and P. D. Seymour. “Bounded-Diameter Tree-Decompositions”. In: *Combinatorica* 44.3 (2024), pp. 659–674. DOI: 10.1007/s00493-024-00088-1. arXiv: 2306.13282.

[B92] D. Bienstock. “Corrigendum: “On the Complexity of Testing for Odd Holes and Induced Odd Paths””. In: *Discrete Mathematics* 102.1 (1992), p. 109. DOI: 10.1016/0012-365X(92)90357-L.

[B91] D. Bienstock. “On the Complexity of Testing for Odd Holes and Induced Odd Paths”. In: *Discrete Mathematics* 90.1 (1991), pp. 85–92. DOI: 10.1016/0012-365X(91)90098-M.

[B93] N. Biggs. *Algebraic Graph Theory*. Cambridge University Press, 1993. 220 pp.

[B48] G. Birkhoff. *Lattice Theory*. revised edition. Colloquium Publications. American Mathematical Society, 1948.

[BP93] J. R. S. Blair and B. Peyton. “An Introduction to Chordal Graphs and Clique Trees”. In: *Graph Theory and Sparse Matrix Computation*. Ed. by A. George, J. R. Gilbert, and J. W. H. Liu. New York, NY: Springer, 1993, pp. 1–29. DOI: 10.1007/978-1-4613-8369-7_1.

[BJKKP] D. Blankenhagen, R. W. Jacobs, P. Knappe, J. Kurkofka, and N. Pich Preuss. “The Local 2-Separator Theorem via Coverings”. In preparation.

[B06] H. L. Bodlaender. “Treewidth: Characterizations, Applications, and Computations”. In: *Graph-Theoretic Concepts in Computer Science*. Ed. by F. V. Fomin. Berlin, Heidelberg: Springer, 2006, pp. 1–14. DOI: 10.1007/11917496_1.

[BGH01] T. Böhme, F. Göring, and J. Harant. “Menger’s Theorem”. In: *Journal of Graph Theory* 37.1 (2001), pp. 35–36. DOI: 10.1002/jgt.1001.

[BZ06] E. D. Bolker and T. Zaslavsky. “A Simple Algorithm That Proves Half-Integrality of Bidirected Network Programming”. In: *Networks* 48.1 (2006), pp. 36–38. DOI: 10.1002/net.20117.

[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”. In: *Journal of the European Mathematical Society* 26.10 (2023), pp. 3739–3791. DOI: 10.4171/jems/1341. arXiv: 2012.02435.

[BDEW24] M. Bonamy, J. Davies, L. Esperet, and A. Wesolek. *Personal Communication*. 2024.

[B21] N. Bowler. *Personal Communication*. 2021.

[BGGJR23] N. Bowler, E. Ghorbani, F. Gut, R. W. Jacobs, and F. Reich. *Menger’s Theorem in Bidirected Graphs*. 2023. arXiv: 2306.01921.

[BGGJR24] N. Bowler, E. Ghorbani, F. Gut, R. W. Jacobs, and F. Reich. *Hitting Cycles through Prescribed Vertices or Edges*. To appear in *SIAM Journal on Discrete Mathematics*. 2024. arXiv: 2412.06557.

[BGHKMR23] N. Bowler, F. Gut, M. Hatzel, K.-i. Kawarabayashi, I. Muzi, and F. Reich. *Decomposition of (Infinite) Digraphs along Directed 1-Separations*. 2023. arXiv: 2305.09192.

[BDS05] H. Bruhn, R. Diestel, and M. Stein. "Menger's Theorem for Infinite Graphs with Ends". In: *Journal of Graph Theory* 50.3 (2005), pp. 199–211. DOI: 10.1002/jgt.20108.

[BS07] H. Bruhn and M. Stein. "On End Degrees and Infinite Cycles in Locally Finite Graphs". In: *Combinatorica* 27.3 (2007), pp. 269–291. DOI: 10.1007/s00493-007-2149-0.

[BFT25] B. M. Bumpus, J. Fairbanks, and W. J. Turner. *Lassos: Pushing Tree Decompositions Forward Along Homomorphisms*. 2025. arXiv: 2408.15184.

[BKMM25] B. M. Bumpus, Z. A. Kocsis, J. E. Master, and E. Minichiello. *Structured Decompositions: Structural and Algorithmic Compositionality*. 2025. arXiv: 2207.06091.

[BK22] C. Bürger and J. Kurkofka. "Duality Theorems for Stars and Combs IV: Undominating Stars". In: *Journal of Graph Theory* 100.1 (2022), pp. 140–162. DOI: 10.1002/jgt.22769. arXiv: 2004.00591.

[B65] J. P. Burling. "On Coloring Problems of Families of Polytopes". PhD thesis. University of Colorado, 1965. 66 pp.

[CDHH14] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. "k-Blocks: A Connectivity Invariant for Graphs". In: *SIAM Journal on Discrete Mathematics* 28.4 (2014), pp. 1876–1891. DOI: 10.1137/130923646. arXiv: 1305.4557.

[CDHH16a] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. "Canonical Tree-Decompositions of Finite Graphs I. Existence and Algorithms". In: *Journal of Combinatorial Theory, Series B* 116 (2016), pp. 1–24. DOI: 10.1016/j.jctb.2014.04.001.

[CDHH16b] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. "Canonical Tree-Decompositions of Finite Graphs II. Essential Parts". In: *Journal of Combinatorial Theory, Series B* 118 (2016), pp. 268–283. DOI: 10.1016/j.jctb.2014.12.009.

[C16] J. Carmesin. "A Short Proof That Every Finite Graph Has a Tree-Decomposition Displaying Its Tangles". In: *European Journal of Combinatorics* 58 (2016), pp. 61–65. DOI: 10.1016/j.ejc.2016.04.007.

[C19] J. Carmesin. "All Graphs Have Tree-Decompositions Displaying Their Topological Ends". In: *Combinatorica* 39.3 (2019), pp. 545–596. DOI: 10.1007/s00493-018-3572-0.

[C22] J. Carmesin. "Local 2-Separators". In: *Journal of Combinatorial Theory, Series B* 156 (2022), pp. 101–144. DOI: 10.1016/j.jctb.2022.04.005. arXiv: 2008.03032.

[C] J. Carmesin. “Local k-Separators”. Unpublished manuscript.

[CDHS14] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. “Connectivity and Tree Structure in Finite Graphs”. In: *Combinatorica* 34.1 (2014), pp. 11–46. DOI: 10.1007/s00493-014-2898-5. arXiv: 1105.1611.

[CF23] J. Carmesin and S. Frenkel. *How to Apply Tree Decomposition Ideas in Large Networks?* 2023. arXiv: 2312.12354.

[CG17] J. Carmesin and J. P. Gollin. “Canonical Tree-Decompositions of a Graph That Display Its k-Blocks”. In: *Journal of Combinatorial Theory, Series B* 122 (2017), pp. 1–20. DOI: 10.1016/j.jctb.2016.05.001.

[CHM22] J. Carmesin, M. Hamann, and B. Miraftab. “Canonical Trees of Tree-Decompositions”. In: *Journal of Combinatorial Theory, Series B* 152 (2022), pp. 1–26. DOI: 10.1016/j.jctb.2021.08.004.

[CKKT25] J. Carmesin, G. Kontogeorgiou, J. Kurkofka, and W. J. Turner. *Towards a Stallings-type Theorem for Finite Groups*. 2025. arXiv: 2403.07776.

[CK23] J. Carmesin and J. Kurkofka. “Canonical Decompositions of 3-Connected Graphs”. In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. 2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS). 2023, pp. 1887–1920. DOI: 10.1109/FOCS57990.2023.00115. arXiv: 2304.00945.

[CK24] J. Carmesin and J. Kurkofka. “Entanglements”. In: *Journal of Combinatorial Theory, Series B* 164 (2024), pp. 17–28. DOI: 10.1016/j.jctb.2023.08.007. arXiv: 2205.11488.

[CK25] J. Carmesin and J. Kurkofka. “Characterising 4-Tangles Through a Connectivity Property”. In: *The Electronic Journal of Combinatorics* 3.P3.26 (2025). DOI: 10.37236/12367. arXiv: 2309.00902v2.

[CDNRV12] V. Chepoi, F. F. Dragan, I. Newman, Y. Rabinovich, and Y. Vaxès. “Constant Approximation Algorithms for Embedding Graph Metrics into Trees and Outerplanar Graphs”. In: *Discrete & Computational Geometry* 47.1 (2012), pp. 187–214. DOI: 10.1007/s00454-011-9386-0. arXiv: 1007.0489.

[CDEHV08] V. Chepoi, F. Dragan, B. Estellon, M. Habib, and Y. Vaxès. “Diameters, Centers, and Approximating Trees of Delta-Hyperbolicgeodesic Spaces and Graphs”. In: *Proceedings of the Twenty-Fourth Annual Symposium on Computational Geometry*. SCG ’08. New York, NY, USA: Association for Computing Machinery, 2008, pp. 59–68. DOI: 10.1145/1377676.1377687.

[C77] B. V. Cherkassky. “A Solution of a Problem on Multicommodity Flows in a Network (English Title), Reshenie Odnoi Zadachi o Mnogoproduktovykh Potokakh v Seti (Russian Title)”. In: *Ekonomika i Matematicheskie Metody* 13.1 (1977), pp. 143–151.

[C90] B. Courcelle. “The Monadic Second-Order Logic of Graphs. I. Recognizable Sets of Finite Graphs”. In: *Information and Computation* 85.1 (1990), pp. 12–75. DOI: 10.1016/0890-5401(90)90043-H.

[DHIM24] J. Davies, R. Hickingbotham, F. Illingworth, and R. McCarty. *Fat Minors Cannot Be Thinned (by Quasi-Isometries)*. To appear in *Analysis and Geometry in Metric Spaces*. 2024. arXiv: 2405.09383.

[DDFMMS14] M. Devos, Z. Dvořák, J. Fox, J. McDonald, B. Mohar, and D. Scheide. “A Minimum Degree Condition Forcing Complete Graph Immersion”. In: *Combinatorica* 34.3 (2014), pp. 279–298. DOI: 10.1007/s00493-014-2806-z.

[DMMS13] M. DeVos, J. McDonald, B. Mohar, and D. Scheide. “A Note on Forbidding Clique Immersions”. In: *The Electronic Journal of Combinatorics* (2013), P55–P55. DOI: 10.37236/2533.

[D94] R. Diestel. “The Depth-First Search Tree Structure of $TK^{\aleph_0}$-Free Graphs”. In: *Journal of Combinatorial Theory, Series B* 61.2 (1994), pp. 260–262. DOI: 10.1006/jctb.1994.1048.

[D88] R. Diestel. “Simplicial Minors and Decompositions of Graphs”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 103.3 (1988), pp. 409–426. DOI: 10.1017/S0305004100065026.

[D91] R. Diestel. “Decomposing Infinite Graphs”. In: *Discrete Math.* 95.1 (1991), pp. 69–89.

[D03] R. Diestel. “The Countable Erdős–Menger Conjecture with Ends”. In: *Journal of Combinatorial Theory, Series B* 87.1 (2003), pp. 145–161. DOI: 10.1016/S0095-8956(02)00032-1.

[D17a] R. Diestel. “Ends and Tangles”. In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 87.2 (2017), pp. 223–244. DOI: 10.1007/s12188-016-0163-0. arXiv: 1510.04050.

[D17b] R. Diestel. *Graph Theory*. 5th ed. Vol. 173. Graduate Texts in Mathematics. Berlin, Heidelberg: Springer, 2017. DOI: 10.1007/978-3-662-53622-3.

[D18a] R. Diestel. “Abstract Separation Systems”. In: *Order* 35.1 (2018), pp. 157–170. DOI: 10.1007/s11083-017-9424-5. arXiv: 1406.3797.

[D18b] R. Diestel. “Tree Sets”. In: *Order* 35.1 (2018), pp. 171–192. DOI: 10.1007/s11083-017-9425-4. arXiv: 1512.03781.

[D19] R. Diestel. *Tangles in the Social Sciences*. 2019. arXiv: 1907.07341.

[D20] R. Diestel. *Personal Communication*. 2020.

[D21] R. Diestel. *Homological Aspects of Oriented Hypergraphs*. 2021. arXiv: 2007.09125.

[D24] R. Diestel. *Tangles: A Structural Approach to Artificial Intelligence in the Empirical Sciences*. Cambridge University Press, 2024. URL: tangles-book.com.

[D25] R. Diestel. *Graph Theory*. 6th ed. Vol. 173. Graduate Texts in Mathematics. Berlin, Heidelberg: Springer, 2025. DOI: 10.1007/978-3-662-70107-2.

[DEET23] R. Diestel, C. Elbracht, J. Erde, and M. Teegen. "Duality and Tangles of Set Separations". In: *Journal of Combinatorics* 15.1 (2023), pp. 1–39. DOI: 10.4310/JOC.2024.v15.n1.a1. arXiv: 2109.08398.

[DEJ24] R. Diestel, C. Elbracht, and R. W. Jacobs. "Point Sets and Functions Inducing Tangles of Set Separations". In: *Journal of Combinatorics* 15.3 (2024), pp. 283–306. DOI: 10.4310/JOC.240907010505. arXiv: 2107.01087.

[DEK] R. Diestel, J. Erde, and K. Krishnareddy. *Tangles: A Novel Way to Classify Texts*. In preparation.

[DEW19] R. Diestel, J. Erde, and D. Weißauer. "Structural Submodularity and Tangles in Abstract Separation Systems". In: *Journal of Combinatorial Theory, Series A* 167 (2019), pp. 155–180. DOI: 10.1016/j.jcta.2019.05.001. arXiv: 1805.01439.

[DHL19] R. Diestel, F. Hundertmark, and S. Lemanczyk. "Profiles of Separations: In Graphs, Matroids, and Beyond". In: *Combinatorica* 39.1 (2019), pp. 37–75. DOI: 10.1007/s00493-017-3595-y. arXiv: 1110.6207.

[DJKK25] R. Diestel, R. W. Jacobs, P. Knappe, and J. Kurkofka. *Canonical Graph Decompositions via Coverings, Extended Version*. 2025. arXiv: 2207.04855v7.

[DJKK26] R. Diestel, R. W. Jacobs, P. Knappe, and J. Kurkofka. *Canonical Graph Decompositions via Coverings*. To appear in *Transactions of the AMS*. 2026. arXiv: 2207.04855.

[DJKW25] R. Diestel, R. W. Jacobs, P. Knappe, and P. Wollan. "A Grid Theorem for Strong Immersions of Walls". In: *Journal of Graph Theory* 110.1 (2025). DOI: 10.1002/jgt.23245. arXiv: 2301.05134.

[DK03] R. Diestel and D. Kühn. "Graph-Theoretical versus Topological Ends of Graphs". In: *Journal of Combinatorial Theory, Series B* 87.1 (2003), pp. 197–206. DOI: 10.1016/S0095-8956(02)00034-5.

[DK05] R. Diestel and D. Kühn. "Graph Minor Hierarchies". In: *Discrete Applied Mathematics*. Structural Decompositions, Width Parameters, and Graph Labelings 145.2 (2005), pp. 167–182. DOI: 10.1016/j.dam.2004.01.010.

[DL01] R. Diestel and I. Leader. "Normal Spanning Trees, Aronszajn Trees and Excluded Minors". In: *Journal of the London Mathematical Society* 63.1 (2001), pp. 16–32. DOI: 10.1112/S0024610700001708.

[DO19] R. Diestel and S.-i. Oum. “Tangle-Tree Duality: In Graphs, Matroids and Beyond”. In: *Combinatorica* 39.4 (2019), pp. 879–910. DOI: 10.1007/s00493-019-3798-5. arXiv: 1701.02651.

[DO21] R. Diestel and S.-i. Oum. “Tangle-Tree Duality in Abstract Separation Systems”. In: *Advances in Mathematics* 377 (2021), p. 107470. DOI: 10.1016/j.aim.2020.107470. arXiv: 1701.02509.

[D61] G. A. Dirac. “On Rigid Circuit Graphs”. In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 25.1 (1961), pp. 71–76. DOI: 10.1007/BF02992776.

[D74] D. Djoković. “Automorphisms of Graphs and Coverings”. In: *Journal of Combinatorial Theory, Series B* 16.3 (1974), pp. 243–247. DOI: 10.1016/0095-8956(74)90070-7.

[DG07] Y. Dourisboure and C. Gavoille. “Tree-Decompositions with Bags of Small Diameter”. In: *Discrete Mathematics*. EuroComb ’03 - Graphs and Algorithms 307.16 (2007), pp. 2008–2029. DOI: 10.1016/j.disc.2005.12.060.

[DK14a] F. F. Dragan and E. Köhler. “An Approximation Algorithm for the Tree T-Spanner Problem on Unweighted Graphs via Generalized Chordal Graphs”. In: *Algorithmica* 69.4 (2014), pp. 884–905. DOI: 10.1007/s00453-013-9765-4.

[D85] M. J. Dunwoody. “The Accessibility of Finitely Presented Groups”. In: *Inventiones mathematicae* 81.3 (1985), pp. 449–457. DOI: 10.1007/BF01388581.

[DK15] M. J. Dunwoody and B. Krön. “Vertex Cuts”. In: *Journal of Graph Theory* 80.2 (2015), pp. 136–171. DOI: 10.1002/jgt.21844.

[DK14b] Z. Dvořák and T. Klimošová. “Strong Immersions and Maximum Degree”. In: *SIAM Journal on Discrete Mathematics* 28.1 (2014), pp. 177–187. DOI: 10.1137/130915467. arXiv: 1304.0728.

[DW16] Z. Dvořák and P. Wollan. “A Structure Theorem for Strong Immersions”. In: *Journal of Graph Theory* 83.2 (2016), pp. 152–163. DOI: 10.1002/jgt.21990.

[E65] J. Edmonds. “Paths, Trees, and Flowers”. In: *Canadian Journal of Mathematics* 17 (1965), pp. 449–467. DOI: 10.4153/CJM-1965-045-4.

[E15] C. Elbracht. “Tangles Determined by Majority Vote”. MA thesis. Universität Hamburg, 2015.

[EK21] C. Elbracht and J. L. Kneip. “A Canonical Tree-of-Tangles Theorem for Structurally Submodular Separation Systems”. In: *Combinatorial Theory* 1 (2021). DOI: 10.5070/C61055362. arXiv: 2009.02091.

[EKT20] C. Elbracht, J. L. Kneip, and M. Teegen. “Tangles Are Decided by Weighted Vertex Sets”. In: *Advances in Combinatorics* 2020:9 (2020). DOI: 10.19086/aic.13691. arXiv: 1811.06821.

[EKT21] C. Elbracht, J. L. Kneip, and M. Teegen. “Trees of Tangles in Abstract Separation Systems”. In: *Journal of Combinatorial Theory, Series A* 180 (2021), p. 105425. DOI: 10.1016/j.jcta.2021.105425. arXiv: 1909.09030.

[EKT22a] C. Elbracht, J. L. Kneip, and M. Teegen. “Obtaining Trees of Tangles from Tangle-Tree Duality”. In: *Journal of Combinatorics* 13.2 (2022), pp. 251–287. DOI: 10.4310/JOC.2022.v13.n2.a3. arXiv: 2011.09758.

[EKT22b] C. Elbracht, J. L. Kneip, and M. Teegen. “Trees of Tangles in Infinite Separation Systems”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 173.2 (2022), pp. 297–327. DOI: 10.1017/S0305004121000530. arXiv: 2005.12122.

[EK22] A.-K. Elm and J. Kurkofka. “A Tree-of-Tangles Theorem for Infinite Tangles”. In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 92.2 (2022), pp. 139–178. DOI: 10.1007/s12188-022-00259-1. arXiv: 2003.02535.

[E17] J. Erde. “Refining a Tree-Decomposition Which Distinguishes Tangles”. In: *SIAM Journal on Discrete Mathematics* 31.3 (2017), pp. 1529–1551. DOI: 10.1137/16M1059539. arXiv: 1512.02499.

[E18] J. Erde. “A Unified Treatment of Linked and Lean Tree-Decompositions”. In: *Journal of Combinatorial Theory, Series B* 130 (2018), pp. 114–143. DOI: 10.1016/j.jctb.2017.12.001.

[EP65] P. Erdős and L. Pósa. “On Independent Circuits Contained in a Graph”. In: *Canadian Journal of Mathematics* 17 (1965), pp. 347–352. DOI: 10.4153/CJM-1965-035-8.

[EGL24] L. Esperet, U. Giocanti, and C. Legrand-Duchesne. “The Structure of Quasi-Transitive Graphs Avoiding a Minor with Applications to the Domino Problem”. In: *Journal of Combinatorial Theory, Series B* 169 (2024), pp. 561–613. DOI: 10.1016/j.jctb.2024.08.002. arXiv: 2304.01823.

[FGTW08] M. Ferrara, R. J. Gould, G. Tansey, and T. Whalen. “On H-immersions”. In: *Journal of Graph Theory* 57.3 (2008), pp. 245–254. DOI: 10.1002/jgt.20283.

[FP23] K. Fujiwara and P. Papasoglu. *A Coarse-Geometry Characterization of Cacti*. 2023. arXiv: 2305.08512.

[GKL23] P. Gartland, T. Korhonen, and D. Lokshtanov. *On Induced Versions of Menger’s Theorem on Sparse Graphs*. 2023. arXiv: 2309.08169.

[G74] F. Gavril. “The Intersection Graphs of Subtrees in Trees Are Exactly the Chordal Graphs”. In: *Journal of Combinatorial Theory, Series B* 16.1 (1974), pp. 47–56. DOI: 10.1016/0095-8956(74)90094-X.

[GGRW06] J. Geelen, B. Gerards, N. Robertson, and W. Whittle Geoff. "Obstructions to Branch-Decomposition of Matroids". In: *Journal of Combinatorial Theory, Series B* 96.4 (2006), pp. 560–570. DOI: 10.1016/j.jctb.2005.11.001.

[GKMW23] J. Geelen, O.-j. Kwon, R. McCarty, and P. Wollan. "The Grid Theorem for Vertex-Minors". In: *Journal of Combinatorial Theory, Series B*. Robin Thomas 1962-2020 158 (2023), pp. 93–116. DOI: 10.1016/j.jctb.2020.08.004. arXiv: 1909.08113.

[G17] A. Georgakopoulos. "On Covers of Graphs by Cayley Graphs". In: *European Journal of Combinatorics* 64 (2017), pp. 57–65. DOI: 10.1016/j.ejc.2017.03.013. arXiv: 1504.00173.

[GP25] A. Georgakopoulos and P. Papasoglu. "Graph Minors and Metric Spaces". In: *Combinatorica* 45.3 (2025), p. 33. DOI: 10.1007/s00493-025-00150-6. arXiv: 2305.07456.

[GKW24] A. C. Giannopoulou, S. Kreutzer, and S. Wiederrecht. "Excluding a Planar Matching Minor in Bipartite Graphs". In: *Journal of Combinatorial Theory, Series B* 164 (2024), pp. 161–221. DOI: 10.1016/j.jctb.2023.09.003. arXiv: 2106.00703.

[GHW25] J. P. Gollin, M. Hatzel, and S. Wiederrecht. *On Graphs with a Simple Structure of Maximal Cliques*. 2025. arXiv: 2504.16863.

[GH22] J. P. Gollin and K. Heuer. "Characterising K-Connected Sets in Infinite Graphs". In: *Journal of Combinatorial Theory, Series B* 157 (2022), pp. 451–499. DOI: 10.1016/j.jctb.2022.08.004. arXiv: 1811.06411.

[GK21] J. P. Gollin and J. L. Kneip. "Representations of Infinite Tree Sets". In: *Order* 38.1 (2021), pp. 79–96. DOI: 10.1007/s11083-020-09529-0. arXiv: 1908.10327.

[G16] M. Grohe. "Tangles and Connectivity in Graphs". In: *Language and Automata Theory and Applications*. Ed. by A.-H. Dediu, J. Janoušek, C. Martín-Vide, and B. Truthe. Cham: Springer International Publishing, 2016, pp. 24–41. DOI: 10.1007/978-3-319-30000-9_2. arXiv: 1602.04727.

[GKMW11] M. Grohe, K.-i. Kawarabayashi, D. Marx, and P. Wollan. "Finding Topological Subgraphs Is Fixed-Parameter Tractable". In: *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*. STOC '11. New York, NY, USA: Association for Computing Machinery, 2011, pp. 479–488. DOI: 10.1145/1993636.1993700. arXiv: 1011.1827.

[GS15a] M. Grohe and P. Schweitzer. "Computing with Tangles". In: *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*. STOC '15. New York, NY, USA: Association for Computing Machinery, 2015, pp. 683–692. DOI: 10.1145/2746539.2746587. 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”. In: *IEEE 56th Annual Symposium on Foundations of Computer Science*. 2015, pp. 1010–1029. DOI: 10.1109/FOCS.2015.66. arXiv: 1505.03737.

[G93] M. Gromov. *Geometric Group Theory, Volume 2: Asymptotic Invariants of Infinite Groups*. Cambridge University Press, 1993. 307 pp.

[H91a] R. Halin. “Lattices of Cuts in Graphs”. In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 61.1 (1991), pp. 217–230. DOI: 10.1007/BF02950767.

[H91b] R. Halin. “Tree-Partitions of Infinite Graphs”. In: *Discrete Mathematics* 97.1–3 (1991), pp. 203–217. DOI: 10.1016/0012-365X(91)90436-6.

[H64] R. Halin. “Über simpliziale Zerfällungen beliebiger (endlicher oder unendlicher) Graphen”. In: *Mathematische Annalen* 156.3 (1964), pp. 216–225. DOI: 10.1007/BF01363288.

[H65] R. Halin. “Über Die Maximalzahl Fremder Unendlicher Wege in Graphen”. In: *Mathematische Nachrichten* 30.1–2 (1965), pp. 63–85. DOI: 10.1002/mana.19650300106.

[H75] R. Halin. “Some Path Problems in Graph Theory”. In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 44.1 (1975), pp. 175–186. DOI: 10.1007/BF02992956.

[H77] R. Halin. “Systeme Disjunkter Unendlicher Wege in Graphen”. In: *Numerische Methoden bei Optimierungsaufgaben Band 3: Optimierung bei graphentheoretischen und ganzzahligen Problemen*. Ed. by L. Collatz, G. Meinardus, and W. Wetterling. Basel: Birkhäuser, 1977, pp. 55–67. DOI: 10.1007/978-3-0348-5936-3_4.

[H78] R. Halin. “Simplicial Decompositions of Infinite Graphs”. In: *Annals of Discrete Mathematics*. Ed. by B. Bollobás. Vol. 3. Advances in Graph Theory. Elsevier, 1978, pp. 93–109. DOI: 10.1016/S0167-5060(08)70500-4.

[H89] R. Halin. *Graphentheorie*. De Gruyter, 1989.

[H18] M. Hamann. “Accessibility in Transitive Graphs”. In: *Combinatorica* 38.4 (2018), pp. 847–859. DOI: 10.1007/s00493-017-3361-1. arXiv: 1404.7677.

[HLMR22] M. Hamann, F. Lehner, B. Miraftab, and T. Rühmann. “A Stallings Type Theorem for Quasi-Transitive Graphs”. In: *Journal of Combinatorial Theory, Series B* 157 (2022), pp. 40–69. DOI: 10.1016/j.jctb.2022.05.008.

[HW17] D. J. Harvey and D. R. Wood. “Parameters Tied to Treewidth”. In: *Journal of Graph Theory* 84.4 (2017), pp. 364–385. DOI: 10.1002/jgt.22030. arXiv: 1312.3401.

[H02] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. 574 pp. URL: https://pi.math.cornell.edu/~hatcher/AT/AT.pdf.

[HRW19] M. Hatzel, R. Rabinovich, and S. Wiederrecht. "Cyclewidth and the Grid Theorem for Perfect Matching Width of Bipartite Graphs". In: *Graph-Theoretic Concepts in Computer Science*. Ed. by I. Sau and D. M. Thilikos. Cham: Springer International Publishing, 2019, pp. 53–65. DOI: 10.1007/978-3-030-30786-8_5. arXiv: 1902.01322.

[HU18] M. Heinlein and A. Ulmer. *Directed Cycles Have the Edge-Erdős-Pósa Property*. 2018. arXiv: 1802.05026.

[HNST24] K. Hendrey, S. Norin, R. Steiner, and J. Turcotte. "On an Induced Version of Menger's Theorem". In: *The Electronic Journal of Combinatorics* (2024), P4.28–P4.28. ISSN: 1077-8926. DOI: 10.37236/12575. eprint: 2309.07905.

[HL18] C. Hofer-Temmel and F. Lehner. "Clique Trees of Infinite Locally Finite Chordal Graphs". In: *The Electronic Journal of Combinatorics* (2018), P2.9–P2.9. DOI: 10.37236/3928. arXiv: 1311.7001.

[JJKKM23] R. W. Jacobs, A. Joó, P. Knappe, J. Kurkofka, and R. Melcher. "The Lovász-Cherkassky Theorem for Locally Finite Graphs with Ends". In: *Discrete Mathematics* 346.12 (2023), p. 113586. DOI: 10.1016/j.disc.2023.113586. arXiv: 2109.11260.

[JK24] R. W. Jacobs and P. Knappe. "Efficiently Distinguishing All Tangles in Locally Finite Graphs". In: *Journal of Combinatorial Theory, Series B* 167 (2024), pp. 189–214. DOI: 10.1016/j.jctb.2024.03.004. arXiv: 2303.09332.

[JK25] R. W. Jacobs and P. Knappe. *Canonical Tree-Decompositions of Chordal Graphs*. 2025. arXiv: 2512.18480.

[JKK25] R. W. Jacobs, P. Knappe, and J. Kurkofka. *Canonical Graph Decompositions via Local Separations*. 2025. arXiv: 2501.16170v2.

[J84] K. Jänich. *Topology*. Undergraduate Texts in Mathematics. Springer New York, 1984. 193 pp. URL: https://link.springer.com/book/9780387908922.

[JRST01] T. Johnson, N. Robertson, P. D. Seymour, and R. Thomas. "Directed Tree-Width". In: *Journal of Combinatorial Theory, Series B* 82.1 (2001), pp. 138–154. DOI: 10.1006/jctb.2000.2031.

[J23] A. Joó. "The Lovász-Cherkassky Theorem in Countable Graphs". In: *Journal of Combinatorial Theory, Series B* 159 (2023), pp. 1–19. DOI: 10.1016/j.jctb.2022.11.001. arXiv: 2102.04203.

[J69] H. A. Jung. "Wurzelbäume Und Unendliche Wege in Graphen". In: *Mathematische Nachrichten* 41.1–3 (1969), pp. 1–22. DOI: 10.1002/mana.19690410102.

[KK12] N. Kakimura and K.-i. Kawarabayashi. “Packing Directed Circuits through Prescribed Vertices Bounded Fractionally”. In: *SIAM Journal on Discrete Mathematics* 26.3 (2012), pp. 1121–1133. DOI: 10.1137/100786423.

[K72] R. M. Karp. “Reducibility among Combinatorial Problems”. In: *Complexity of Computer Computations: Proceedings of a Symposium on the Complexity of Computer Computations, Held March 20–22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, and Sponsored by the Office of Naval Research, Mathematics Program, IBM World Trade Corporation, and the IBM Research Mathematical Sciences Department*. Ed. by R. E. Miller, J. W. Thatcher, and J. D. Bohlinger. The IBM Research Symposia Series. Boston, MA: Springer US, 1972, pp. 85–103. DOI: 10.1007/978-1-4684-2001-2_9.

[KKKK13] K.-i. Kawarabayashi, D. Král’, M. Krčál, and S. Kreutzer. “Packing Directed Cycles through a Specified Vertex Set”. In: *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2013, pp. 365–377. DOI: 10.1137/1.9781611973105.27.

[KK15] K.-i. Kawarabayashi and S. Kreutzer. “The Directed Grid Theorem”. In: *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*. STOC ’15. New York, NY, USA: Association for Computing Machinery, 2015, pp. 655–664. DOI: 10.1145/2746539.2746586. arXiv: 1411.5681.

[KNTW19] K.-i. Kawarabayashi, S. Norine, R. Thomas, and P. Wollan. “K_6 Minors in 6-Connected Graphs of Bounded Tree-Width”. In: *Journal of Combinatorial Theory, Series B* 136 (2019), pp. 1–32. DOI: 10.1016/j.jctb.2017.08.006. arXiv: 1203.2171.

[KEFKRTL23] S. Klepper, C. Elbracht, D. Fioravanti, J. L. Kneip, L. Rendsburg, M. Teegen, and U. von Luxburg. “Clustering with Tangles: Algorithmic Framework and Theoretical Guarantees”. In: *Journal of Machine Learning Research* 24.190 (2023), pp. 1–56. arXiv: 2006.14444. URL: http://jmlr.org/papers/v24/21-1160.html.

[KK24] P. Knappe and J. Kurkofka. “The Immersion-Minimal Infinitely Edge-Connected Graph”. In: *Journal of Combinatorial Theory, Series B* 164 (2024), pp. 492–516. DOI: 10.1016/j.jctb.2023.10.007. arXiv: 2207.08459.

[KT] P. Knappe and M. Trinh. “Fundamental Groups of Planar Graphs”. In preparation.

[KS26] R. Köhl and M. R. Salarian. *Combinatorial Characterizations of Virtually Torsion-Free and Virtually Free Groups*. 2026. arXiv: 2603.04612.

[KKP23] M. Koloschin, T. Krill, and M. Pitz. “End Spaces and Tree-Decompositions”. In: *Journal of Combinatorial Theory, Series B* 161 (2023), pp. 147–179. DOI: 10.1016/j.jctb.2023.02.007. arXiv: 2205.09865.

[K95] P. Komjáth. “A Note on Minors of Uncountable Graphs”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 117.1 (1995), pp. 7–9. DOI: 10.1017/S0305004100072881.

[K36] D. König. *Theorie Der Endlichen Und Unendlichen Graphen*. Akademische Verlagsgesellschaft, 1936.

[K25] T. Korhonen. “Linear-Time Algorithms for k-Edge-Connected Components, k-Lean Tree Decompositions, and More”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. STOC ’25. New York, NY, USA: Association for Computing Machinery, 2025, pp. 111–119. DOI: 10.1145/3717823.3718123. arXiv: 2411.02658.

[K59a] A. Kotzig. “Z teórie konečných grafov s lineárnym faktorom. I.” In: *Matematicko-fyzikálny časopis* 09.2 (1959), pp. 73–91. URL: http://eudml.org/doc/29879.

[K59b] A. Kotzig. “Z teórie konečných grafov s lineárnym faktorom. II.” In: *Matematicko-fyzikálny časopis* 09.3 (1959), pp. 136–159. URL: http://eudml.org/doc/29908.

[K60] A. Kotzig. “Z teórie konečných grafov s lineárnym faktorom. III.” In: *Matematicko-fyzikálny časopis* 10.4 (1960), pp. 205–215. URL: http://eudml.org/doc/29784.

[KLMPW03] D. Kratsch, H.-O. Le, H. Müller, E. Prisner, and D. Wagner. “Additive Tree Spanners”. In: *SIAM Journal on Discrete Mathematics* 17.2 (2003), pp. 332–340. DOI: 10.1137/S0895480195295471.

[KT91] I. Kříž and R. Thomas. “The Menger-like Property of the Tree-Width of Infinite Graphs”. In: *Journal of Combinatorial Theory, Series B* 52.1 (1991), pp. 86–91. DOI: 10.1016/0095-8956(91)90093-Y.

[K10] B. Krön. “Cutting up Graphs Revisited – a Short Proof of Stallings’ Structure Theorem”. In: 2.2 (2010), pp. 213–221. DOI: 10.1515/gcc.2010.013. arXiv: 1003.1096.

[KMP22] J. Kurkofka, R. Melcher, and M. Pitz. “A Strengthening of Halin’s Grid Theorem”. In: *Mathematika* 68.4 (2022), pp. 1009–1013. DOI: 10.1112/mtk.12135. arXiv: 2104.10672.

[KP26] J. Kurkofka and T. Planken. “A Tutte-type Canonical Decomposition of 3- and 4-Connected Graphs”. In: *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Proceedings. Society for Industrial and Applied Mathematics, 2026, pp. 2942–3021. DOI: 10.1137/1.9781611978971.110. eprint: 2504.00760.

[L76] László Lovász. “On Some Connectivity Properties of Eulerian Graphs”. In: *Acta Mathematica Academiae Scientiarum Hungaricae* 28.1–2 (1976), pp. 129–138. DOI: 10.1007/bf01902503.

[L17a] A. Leitert. “Tree-Breadth of Graphs with Variants and Applications”. PhD thesis. Kent State University, 2017. URL: https://etd.ohiolink.edu/acprod/odb_etd/etd/r/1501/10?clear=10&p10_accession_num=kent1497402176814598.

[L17b] C. Löh. *Geometric Group Theory*. Universitext. Cham: Springer International Publishing, 2017. DOI: 10.1007/978-3-319-72254-2.

[MKS04] W. Magnus, A. Karrass, and D. Solitar. *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*. Courier Corporation, 2004. 466 pp.

[M05] J. F. Manning. “Geometry of Pseudocharacters”. In: *Geometry & Topology* 9.2 (2005), pp. 1147–1185. DOI: 10.2140/gt.2005.9.1147. arXiv: math/0303380.

[MW14] D. Marx and P. Wollan. “Immersions in Highly Edge Connected Graphs”. In: *SIAM Journal on Discrete Mathematics* 28.1 (2014), pp. 503–520. DOI: 10.1137/130924056. arXiv: 1305.1331.

[M27] K. Menger. “Zur Allgemeinen Kurventheorie”. In: *Fundamenta Mathematicae* 10 (1927), pp. 96–115. DOI: 10.4064/fm-10-1-96-115.

[M68] J. Milnor. “A Note on Curvature and Fundamental Group”. In: *Journal of Differential geometry* 2.1 (1968), pp. 1–7. DOI: 10.4310/jdg/1214501132.

[NSS25a] T. Nguyen, A. Scott, and P. D. Seymour. “A Counterexample to the Coarse Menger Conjecture”. In: *Journal of Combinatorial Theory, Series B* 173 (2025), pp. 68–82. DOI: 10.1016/j.jctb.2025.01.004.

[NSS25b] T. Nguyen, A. Scott, and P. D. Seymour. *Asymptotic Structure. II. Path-width and Additive Quasi-Isometry*. 2025. arXiv: 2509.09031.

[NSS25c] T. Nguyen, A. Scott, and P. D. Seymour. *Asymptotic Structure. IV. A Counterexample to the Weak Coarse Menger Conjecture*. 2025. arXiv: 2508.14332.

[NSS25d] T. Nguyen, A. Scott, and P. D. Seymour. *Asymptotic Structure. V. The Coarse Menger Conjecture in Bounded Path-Width*. 2025. arXiv: 2509.08762.

[NSS25e] T. Nguyen, A. Scott, and P. D. Seymour. *Asymptotic Structure. VI. Distant Paths across a Disc*. 2025. arXiv: 2509.07174.

[N05] S. Norine. “Matching Structure and Pfaffian Orientation of Graphs”. PhD thesis. Georgia Institute of Technology, 2005. 149 pp.

[P68] H. Perfect. “Applications of Menger’s Graph Theorem”. In: *Journal of Mathematical Analysis and Applications* 22.1 (1968), pp. 96–111. DOI: 10.1016/0022-247X(68)90163-7.

[P17] C.-t. Perng. “On a Class of Theorems Equivalent to Farkas’s Lemma”. In: *Applied Mathematical Sciences* 11 (2017), pp. 2175–2184. DOI: 10.12988/ams.2017.77218.

[P20] M. Pitz. “A Unified Existence Theorem for Normal Spanning Trees”. In: *Journal of Combinatorial Theory, Series B* 145 (2020), pp. 466–469. DOI: 10.1016/j.jctb.2020.07.002. arXiv: 2003.11575.

[P21] M. Pitz. “Proof of Halin’s Normal Spanning Tree Conjecture”. In: *Israel Journal of Mathematics* 246.1 (2021), pp. 353–370. DOI: 10.1007/s11856-021-2249-3. arXiv: 2005.02833.

[P23] M. Pitz. “A Note on Minor Antichains of Uncountable Graphs”. In: *Journal of Graph Theory* 102.3 (2023), pp. 552–555. DOI: 10.1002/jgt.22886. arXiv: 2005.05816.

[P96] N. Polat. “Ends and Multi-Endings, I”. In: *J. Comb. Theory Ser. B* 67.1 (1996), pp. 86–110. DOI: 10.1006/jctb.1996.0035.

[PW12] M. Pontecorvi and P. Wollan. “Disjoint Cycles Intersecting a Set of Vertices”. In: *Journal of Combinatorial Theory, Series B* 102.5 (2012), pp. 1134–1141. DOI: 10.1016/j.jctb.2012.05.004.

[PT23] P. Pournajafi and N. Trotignon. “Burling Graphs Revisited, Part I: New Characterizations”. In: *European Journal of Combinatorics* 110 (2023), p. 103686. DOI: 10.1016/j.ejc.2023.103686. arXiv: 2104.07001.

[PT24] P. Pournajafi and N. Trotignon. “Burling Graphs Revisited, Part II: Structure”. In: *European Journal of Combinatorics* 116 (2024), p. 103849. DOI: 10.1016/j.ejc.2023.103849. arXiv: 2106.16089.

[RRST96] B. Reed, N. Robertson, P. D. Seymour, and R. Thomas. “Packing Directed Circuits”. In: *Combinatorica* 16.4 (1996), pp. 535–554. DOI: 10.1007/BF01271272.

[RS83–04] N. Robertson and P. D. Seymour. “Graph Minors I–XX”. In: *Journal of Combinatorial Theory, Series B* (1983–2004). DOI: 10.1016/0095-8956(83)90079-5.

[RS86a] N. Robertson and P. D. Seymour. “Graph Minors. II. Algorithmic Aspects of Tree-Width”. In: *J.Algorithms* 7 (1986), pp. 309–322. DOI: 10.1016/0196-6774(86)90023-4.

[RS86b] N. Robertson and P. D. Seymour. “Graph Minors. V. Excluding a Planar Graph”. In: *Journal of Combinatorial Theory, Series B* 41 (1986), pp. 92–114. DOI: 10.1016/0095-8956(86)90030-4.

[RS90a] N. Robertson and P. D. Seymour. “Graph Minors. IV. Tree-width and Well-Quasi-Ordering”. In: *J. Combin.Theory (Series B)* 48 (1990), pp. 227–254. DOI: 10.1016/0095-8956(90)90120-O.

[RS90b] N. Robertson and P. D. Seymour. “Graph Minors. VIII. A Kuratowski Theorem for General Surfaces”. In: *J. Combin.Theory (Series B)* 48 (1990), pp. 255–288. DOI: 10.1016/0095-8956(90)90121-F.

[RS91] N. Robertson and P. D. Seymour. “Graph Minors. X. Obstructions to Tree-Decomposition”. In: *Journal of Combinatorial Theory, Series B* 52 (1991), pp. 153–190. DOI: 10.1016/0095-8956(91)90061-N.

[RS95] N. Robertson and P. D. Seymour. “Graph Minors. XIII. The Disjoint Paths Problem”. In: *Journal of Combinatorial Theory, Series B* 63 (1995), pp. 65–110. DOI: 10.1006/jctb.1995.1006.

[RS03] N. Robertson and P. D. Seymour. “Graph Minors. XVI. Excluding a Non-Planar Graph”. In: *J. Combin.Theory (Series B)* 89 (2003), pp. 43–76. DOI: 10.1016/S0095-8956(03)00042-X.

[RS04] N. Robertson and P. D. Seymour. “Graph Minors. XX. Wagner’s Conjecture”. In: *J. Combin.Theory (Series B)* 92 (2004), pp. 325–357. DOI: 10.1016/j.jctb.2004.08.001.

[RS10] N. Robertson and P. D. Seymour. “Graph Minors. XXIII. Nash-Williams’ Immersion Conjecture”. In: *Journal of Combinatorial Theory, Series B* 100 (2010), pp. 181–205. DOI: 10.1016/j.jctb.2009.07.003.

[RST95] N. Robertson, P. D. Seymour, and R. Thomas. “Excluding Infinite Clique Minors”. In: *Memoirs of the American Mathematical Society* 118.566 (1995), pp. vi+103. DOI: 10.1090/memo/0566.

[R17] T. Rühmann. “A Study of Infinite Graphs of a Certain Symmetry and Their Ends”. Doctoral Thesis. Staats- und Universitätsbibliothek Hamburg Carl von Ossietzky, 2017. URL: https://ediss.sub.uni-hamburg.de/handle/ediss/7513.

[S98] A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, 1998. 488 pp.

[S22] A. Scott. “Graphs of Large Chromatic Number”. In: *Proceedings of the International Congress of Mathematics 2022*. 2022, pp. 4660–4680. DOI: 10.4171/icm2022/149.

[ST93] P. D. Seymour and R. Thomas. “Excluding Infinite Trees”. In: *Transactions of the American Mathematical Society* 335.2 (1993), pp. 597–630. DOI: 10.1090/S0002-9947-1993-1079058-6.

[S83] J. R. Stallings. “Topology of Finite Graphs”. In: *Inventiones mathematicae* 71.3 (1983), pp. 551–565. DOI: 10.1007/BF02095993.

[S72] J. R. Stallings. *Group Theory and Three-Dimensional Manifolds*. Yale University Press, 1972.

[Š55] A. S. Švarc. “A Volume Invariant of Coverings”. In: *Dokl. Akad. Nauk SSSR (NS)*. Vol. 105. 1955, pp. 32–34.

[T88] R. Thomas. “A Counter-Example to ‘Wagner’s Conjecture’ for Infinite Graphs”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 103.1 (1988), pp. 55–57. DOI: 10.1017/S0305004100064616.

[T89] R. Thomas. “Well-Quasi-Ordering Infinite Graphs with Forbidden Finite Planar Minor”. In: *Transactions of the American Mathematical Society* 312.1 (1989), pp. 279–313. DOI: 10.1090/S0002-9947-1989-0932450-9.

[T90] R. Thomas. “A Menger-like Property of Tree-Width: The Finite Case”. In: *Journal of Combinatorial Theory, Series B* 48.1 (1990), pp. 67–76. DOI: 10.1016/0095-8956(90)90130-R.

[TW93] C. Thomassen and W. Woess. “Vertex-Transitive Graphs and Accessibility”. In: *Journal of Combinatorial Theory, Series B* 58.2 (1993), pp. 248–268. DOI: 10.1006/jctb.1993.1042.

[T84] W. T. Tutte. *Graph Theory*. Vol. 21. Encyclopedia of Mathematics and Its Applications. Addison-Wesley Publishing Company, Advanced Book Program, Reading, MA, 1984.

[W37] K. Wagner. “Über eine Eigenschaft der ebenen Komplexe”. In: *Mathematische Annalen* 114.1 (1937), pp. 570–590. DOI: 10.1007/BF01594196.

[W71] C. T. C. Wall. “Pairs of Relative Cohomological Dimension One”. In: *Journal of Pure and Applied Algebra* 1.2 (1971), pp. 141–154. DOI: 10.1016/0022-4049(71)90014-4.

[W22] S. Wiederrecht. “Matching Minors in Bipartite Graphs”. Universitätsverlag der Technischen Universität Berlin, 2022. DOI: 10.14279/depositonce-14958.

[W15] P. Wollan. “The Structure of Graphs Not Admitting a Fixed Immersion”. In: *Journal of Combinatorial Theory, Series B* 110 (2015), pp. 47–66. DOI: 10.1016/j.jctb.2014.07.003. arXiv: 1302.3867.

Appendix

APPENDIX A

Summary

Tangles are an abstract notion of highly connected substructures or 'clusters' in graphs. They capture clusters indirectly, by a system of pointers towards the cluster's location in the graph. The central two tangle theorems for finite graphs assert that the tangles of a graph are arranged in a tree-like way (*Tree-of-Tangles Theorem*) and that the absence of a high-order tangle (that is, a large cluster) is witnessed by a tree-decomposition into small parts, each too small to accommodate a tangle of high order (*Tangle-Tree Duality Theorem*). In the last two decades, tangles were axiomatised in a more formal, abstract setting, which allows these methods and results to be applied in other discrete contexts such as matroids or data sets. In this way, tangles provide a unifying framework for the study of various kinds of connectivity and clusters in graphs and beyond.

In this thesis, we study different types of connectivity in graphs through the lens of tangles and separations. We exhibit the global and local aspects of a graph's structure, study the connections of tangles with concrete substructures in a graph, find a tree-decomposition that displays the infinities of an infinite graph, and investigate the interplay of graph-theoretic (connectivity) notions with the large-scale geometry of a graph.

In Part I, we use tangle-theory to display the connectivity structure of a graph, building on the novel concept of 'r-local coverings'. Given an integer parameter $r > 0$ describing the intended degree of locality, the *r-local covering* is a topological covering of a given graph (viewed as a 1-complex) that reflects the graph's r-local structure (that is, its $r/2$-balls), but 'unfolds' every further global structure in a tree-like way. With local coverings we offer a possible formal basis for studying the local and global aspects of graph properties: by studying their interplay with the r-local coverings of the graph. In Part I of this thesis, we formally develop this novel approach and use it to investigate the connectivity structure of graphs by applying concepts from tangle-theory.

In Chapter 1, we use this approach to capture the r-global structure of a graph G by the means of a suitable graph-decomposition $\mathcal{H}_r(G)$. By the definition of the r-local cover G_r, its global structure is tree-like, and we can thus capture it by applying a suitable Tree-of-Tangles Theorem. For the construction of $\mathcal{H}_r(G)$, we then 'fold' the resulting canonical tree-decomposition of G_r along the r-local covering map and obtain the graph-decomposition $\mathcal{H}_r(G)$. Both the decomposition graph and the parts of $\mathcal{H}_r(G)$ are determined by the construction of $\mathcal{H}_r(G)$, not imposed, and they only depend on the chosen integer parameter $r > 0$. In this way, the tangle theory from graph minors is brought to bear canonically on arbitrary graphs, which need not be tree-like.

In Chapter 2, we develop a theory of r-local separations compatible with r-local

coverings and apply it to compute $\mathcal{H}_r(G)$ for many finite graphs G. In contrast to the construction of $\mathcal{H}_r(G)$ in Chapter 1 which relies on the usually infinite r-local covering G_r, we reconstruct the decomposition $\mathcal{H}_r(G)$ just from finite information in G itself that is sufficiently local to be reflected in the cover. Our construction in particular enables algorithmic applications of $\mathcal{H}_r(G)$.

In Chapter 3, we study canonical tree-decompositions of chordal graphs and apply them to characterise r-locally chordal graphs in terms of their $\mathcal{H}_r(G)$. We show that a graph G is r-locally chordal (that is, its $r/2$-balls are chordal) if and only if the graph-decomposition $\mathcal{H}_r(G)$ is into cliques (that is, its parts are complete subgraphs of G).

In Part II, we investigate the relationship between tangles and concrete substructures (such as clusters) in graphs, contributing in particular to the problem of finding vertex sets that induce a given tangle. A cluster of a graph induces a tangle by orienting each low-order separation towards the side that contains most of the cluster. This raises the question whether the converse also holds, that is, whether every tangle in a graph is induced by some cluster via majority vote; the following weak form of this question was proposed by Diestel, Hundertmark and Lemanczyk: *Is every k-tangle τ in a graph induced by some vertex set X by majority vote in that $|X \cap A| < |X \cap B|$ for every separation $(A, B) \in \tau$?*

In Chapter 4, we reduce this question to graphs of size bounded in k, using a new inductive proof method for k-tangles. We also show that if this question has a positive answer for some fixed k, then every k-tangle is induced by a vertex set whose size is bounded in k. Analogously, building on the existence of tangle-inducing weighted vertex sets which was established by Elbracht, Kneip and Teegen, we show for all k that every k-tangle in a graph G is induced by a weight function $w : V(G) \to \mathbb{Z}_{\geqslant 0}$ whose total weight $w(V(G))$ is bounded in k.

In Chapter 5, we study tangles in discrete contexts more general than graphs and provide two characterisations for the existence of point sets X inducing such tangles. The first is in terms of how many small sides of a tangle's separations it takes to cover the ground set. The second uses a new notion of duality for oriented set separations that becomes possible if these are no longer required to be separations of a graph or matroid.

In Chapter 6, we find a concrete substructure of a graph not from a single tangle, but from the graph's tangle-structure: We show that Robertson and Seymour's Tree-of-Tangles Theorem extends to locally finite graphs without thick ends, by exhibiting thick ends (and hence half-grid minors) at every obstruction to the existence of such tangle-distinguishing tree-decompositions.

In Part III, we study how tree-decompositions interact with various notions of connectivity in graphs, especially constructing a tree-decomposition that displays the different types of infinities in infinite graphs and extending Menger's Theorem to bidirected graphs. A recurring topic are structural dualities in terms of width-measures and their respective obstructions, such as the Tangle-Tree Duality Theorem and the various grid theorems.

In Chapter 7, 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. In Chapter 8, we argue that this result achieves very much the best one can hope for, by

providing counterexamples to various possible generalisations.

As a first application of the main result of Chapter 7, 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 its ends and critical vertex sets, strengthening results by Carmesin and by Elm and Kurkofka for this graph class.

As a second application, we show that every graph G which admits a tree-decomposition into finite parts has a tree-decomposition into finite parts that displays all infinities of G, that is, it displays the ends of G homeomorphically, their combined degree and their dominating vertices (resolving a question of Halin from 1977) as well as the critical vertex sets of G and their tight components cofinitely. This latter tree-decomposition yields unified proofs of the characterisations of graphs without half-grid minor and of graphs without binary tree subdivision, which are due to Robertson, Seymour and Thomas. Moreover, Albrechtsen applied this tree-decomposition to extend the Tangle-Tree Duality Theorem of Robertson and Seymour to infinite graphs.

In Chapter 9, we prove an exact structural duality for strong immersions of walls in terms of special tree-cut decompositions. For the case of walls, this improves Dvořák and Wollan's structure theorem for graphs excluding a strong immersion of a fixed graph; this theorem admits no converse for walls.

In Chapter 10, we extend the Lovász-Cherkassky theorem from locally finite graphs to locally finite graphs with ends. For a set T of vertices of a finite graph, the Lovász-Cherkassky theorem studies edge-disjoint T-paths and shows that a natural upper bound is attained if all vertices outside T have even degree. Jóo generalised this to locally finite graphs, and we further extend it to the setting where T may contain both vertices and ends, with paths becoming infinite when connecting to an end.

In Chapter 11, we investigate whether and how Menger's theorem extends to bidirected graphs (a generalisation of (un)directed graphs where every edge has a sign $+$ or $-$ at each of its endvertices). We show that Menger's theorem is in general not true in bidirected graphs. Nonetheless, we introduce a sufficient condition for X and Y which yields a version of Menger's Theorem in bidirected graphs that in particular extends its (un)directed analogues.

In Chapter 12, we prove that for every set S of vertices of a directed graph D, the maximum number of vertices in S contained in a collection of vertex-disjoint cycles in D is at least the minimum size of a set of vertices that hits all cycles containing a vertex of S. We also show that the corresponding statement in bidirected graphs is true and that its edge-variant holds in both undirected and directed graphs, but fails in bidirected graphs; the vertex-version in undirected graphs remains an open problem. As an application, we derive that the (bi)directed tree-width of a (bi)directed graph is linearly bounded in its cycle-width, which improves the previously known quadratic upper bound.

In Part IV, we look at the large-scale geometry of graphs through tools and methods from connectivity theory. The term 'coarse graph theory' was introduced by Georgakopoulos and Papasoglu, who followed Gromov's approach to coarse geometry into the realm of graphs

and presented a collection of problems and results on the interplay of graph minors and geometry. As for Gromov, quasi-isometries also become the central tool of Georgakopoulos' and Papasoglu's approach: They provide the formal framework to study the coarse geometry of graphs.

In Chapters 13 and 14, we prove several cases of the following conjecture of Georgakopoulos and Papasoglu on 'fat minors', a metric variant of minors: *Let $\mathcal{X}$ be a finite set of finite graphs. Then there exists a function $f : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ such that, for all $K \in \mathbb{N}$, every graph with no K-fat X minor for any $X \in \mathcal{X}$ is $f(K)$-quasi-isometric to a graph with no X minor for any $X \in \mathcal{X}$.* Our results build on a novel perspective on quasi-isometries in terms of graph-decompositions: We show that quasi-isometries with bounded parameters can equivalently be described by honest graph-decompositions with bounded radial width and radial spread.

In Chapter 13, we prove this conjecture for $\mathcal{X} = \{K^3, K_{1,3}\}$ (quasi-isometric to a disjoint union of paths), $\mathcal{X} = \{K_{1,3}\}$ (cycles and paths) and $\mathcal{X} = \{K^3, W\}$ (subdivided stars), where W is the 'wrench graph' obtained from two paths of length 2 by joining their middle vertices with an edge. In all these cases, our results even hold true when replacing 'fat minors' by the stronger notion of 'quasi-geodesic topological minors', which can be thought of as a coarse analogue of topological minors.

In Chapter 14, we confirm the above 'fat minor'-conjecture for $\mathcal{X} = \{K^4\}$. Our proof technique also yields a new short proof of the respective K_4^--case, which was first established by Fujiwara and Papasoglu.

In Chapter 15, we prove an approximate metric Menger-type theorem for two paths: We show that for all $d \in \mathbb{N}$ and for all graphs G and vertex sets $X, Y \subseteq V(G)$, either there exist 2 disjoint X–Y paths, pairwise at distance at least d, or there exists a vertex $v \in V(G)$ such that each X–Y path has distance at most $129d$ to v. This result also serves as an important tool in Chapter 14.

APPENDIX B

Deutschsprachige Zusammenfassung

Tangles sind ein abstrakter Begriff für hoch zusammenhängende Teilstrukturen oder „Cluster“ in Graphen. Sie erfassen Cluster indirekt durch ein System von Wegweisern, die auf die Position des Clusters im Graphen zeigen. Die beiden zentralen Tangle-Sätze für endliche Graphen besagen, dass die Tangles eines Graphen baumartig angeordnet sind (*Tree-of-Tangles Theorem*) und dass die Abwesenheit eines Tangles hoher Ordnung (also eines großen Clusters) durch eine Baumzerlegung in kleine Teile belegt wird, von denen jeder zu klein ist, um ein Tangle hoher Ordnung zu enthalten (*Tangle-Tree Duality Theorem*). In den letzten zwei Jahrzehnten wurden Tangles in einem formaleren abstrakten Rahmen axiomatisiert, was eine Anwendung dieser Methoden und Resultate auch in anderen diskreten Kontexten wie Matroiden oder Datensätzen ermöglicht. Auf diese Weise bieten Tangles ein einheitliches Framework für die Untersuchung verschiedener Typen von Zusammenhang und Clustern in Graphen und darüber hinaus.

In dieser Arbeit untersuchen wir verschiedene Arten von Zusammenhang in Graphen aus der Perspektive von Tangles und Teilungen. Wir zeigen die globalen und lokalen Aspekte der Struktur eines Graphen auf, untersuchen die Verbindungen von Tangles mit konkreten Teilstrukturen in einem Graphen, finden eine Baumzerlegung, die die verschiedenen Unendlichkeiten eines unendlichen Graphen darstellt, und erforschen das Zusammenspiel von graphentheoretischen (Zusammenhangs-)Begriffen mit der groben Geometrie eines Graphen.

In Teil I nutzen wir Tangles, um die Zusammenhangsstruktur eines Graphen darzustellen; wir bauen dabei auf dem neu entwickelten Konzept von ‘r-lokalen Überlagerungen’ auf. Für einen gegebenen ganzzahligen Parameter $r > 0$, der den gewünschten Grad der Lokalität festlegt, ist die *r-lokale Überlagerung* eine topologische Überlagerung eines gegebenen Graphen (betrachtet als 1-Komplex), die die r-lokale Struktur des Graphen (also seine $r/2$-Bälle) widerspiegelt, aber alle darüber hinausgehende globale Struktur baumartig „auffaltet“. Mit lokalen Überlagerungen bieten wir eine mögliche formale Grundlage für die Untersuchung der lokalen und globalen Aspekte von Grapheninvarianten: durch die Untersuchung ihres Zusammenspiels mit den r-lokalen Überlagerungen des Graphen. In Teil I der vorliegenden Arbeit entwickeln wir diesen neuen Ansatz formal und nutzen ihn, um die Zusammenhangsstruktur von Graphen unter Anwendung von Konzepten aus der Tangle-Theorie zu untersuchen.

In Kapitel 1 nutzen wir diesen Ansatz, um die r-globale Struktur eines Graphen G durch eine geeignete Graphenzerlegung $\mathcal{H}_r(G)$ zu erfassen. Nach der Definition der r-lokalen Überlagerung G_r ist ihre globale Struktur baumartig, und wir können diese daher durch Anwendung eines geeigneten Tree-of-Tangles Theorems erfassen. Für die Konstruktion von

$\mathcal{H}_r(G)$ „falten“ wir dann die resultierende kanonische Baumzerlegung von G_r entlang der r-lokalen Überlagerungsabbildung und erhalten die Graphenzerlegung $\mathcal{H}_r(G)$. Sowohl der Zerlegungsgraph als auch die Teile von $\mathcal{H}_r(G)$ werden durch die Konstruktion von $\mathcal{H}_r(G)$ gefunden, sie sind also nicht vorgegeben und hängen nur vom gewählten ganzzahligen Parameter $r > 0$ ab. Auf diese Weise lässt sich die Tangle-Theorie auch kanonisch auf Graphen anwenden, die nicht baumartig sind.

In Kapitel 2 entwickeln wir eine Theorie r-lokaler Teilungen, die mit den r-lokalen Überlagerungen kompatibel ist, und verwenden sie, um $\mathcal{H}_r(G)$ für viele endliche Graphen G zu berechnen. Im Gegensatz zur Konstruktion von $\mathcal{H}_r(G)$ in Kapitel 1, die auf dem meist unendlichen r-lokalen Überdeckungsgraphen G_r beruht, rekonstruieren wir die Zerlegung $\mathcal{H}_r(G)$ von G nur aus endlichen Informationen in G selbst, die lokal genug sind, um sich in der Überlagerung wiederzufinden. Unsere Konstruktion ermöglicht daher insbesondere algorithmische Anwendungen von $\mathcal{H}_r(G)$.

In Kapitel 3 untersuchen wir kanonische Baumzerlegungen von chordalen Graphen und verwenden sie, um r-lokal chordale Graphen anhand ihrer $\mathcal{H}_r(G)$ zu charakterisieren. Wir zeigen, dass ein Graph G genau dann r-lokal chordal ist (d. h. seine $r/2$-Kugeln sind chordal), wenn $\mathcal{H}_r(G)$ eine Graphenzerlegung in Cliquen ist (d. h. ihre Teile sind vollständige Teilgraphen von G).

In Teil II untersuchen wir die Beziehung zwischen Tangles und konkreten Teilstrukturen (wie Clustern) in Graphen und tragen dabei insbesondere zur Lösung des Problems bei, Eckenmengen zu finden, die ein gegebenes Tangle induzieren. Ein Cluster eines Graphen induziert ein Tangle, indem jede Teilung niedriger Ordnung in Richtung der Seite orientiert wird, die den größten Teil des Clusters enthält. Dies wirft die Frage auf, ob auch die umgekehrte Implikation gilt, ob also jedes Tangle in einem Graphen durch einen Cluster mittels Mehrheitsbeschluss induziert wird. Die folgende schwache Form dieser Frage wurde von Diestel, Hundertmark und Lemanczyk vorgeschlagen: *Wird jedes k-Tangle τ in einem Graphen durch eine Eckenmenge X per Mehrheitsbeschluss induziert, indem $|X \cap A| < |X \cap B|$ für jede Teilung $(A, B) \in \tau$ gilt?*

In Kapitel 4 reduzieren wir diese Frage auf Graphen, deren Größe in k beschränkt ist; dabei wenden wir eine neue induktive Beweismethode für k-Tangles an. Falls diese Frage für ein gegebenes k eine positive Antwort hat, zeigen wir außerdem, dass jedes k-Tangle durch eine Eckenmenge induziert wird, deren Größe in k beschränkt ist. Für alle k zeigen wir analog dazu, dass jedes k-Tangle in einem Graphen G durch eine Gewichtungsfunktion $w : V(G) \to \mathbb{Z}_{\geqslant 0}$ induziert wird, deren Gesamtgewicht $w(V(G))$ in k beschränkt ist; dies baut auf einem Resultat von Elbracht, Kneip und Teegen auf, welches die Existenz von Tangle-induzierenden gewichteten Eckenmengen zeigt.

In Kapitel 5 untersuchen wir Tangles in allgemeineren diskreten Kontexten als Graphen und zeigen zwei Charakterisierungen für die Existenz von Punktmengen X, die solche Tangles induzieren. Die erste Charakterisierung bezieht sich darauf, wie viele kleine Seiten der Teilungen eines Tangles benötigt werden, um die gesamte Grundmenge zu überdecken. Die zweite Charakterisierung basiert auf einem neuen Dualitätsbegriff für orientierte Mengenteilungen, der dann möglich wird, wenn diese nicht mehr Teilungen eines Graphen oder Matroids sein müssen.

In Kapitel 6 finden wir eine konkrete Teilstruktur eines Graphen nicht durch einen einzelnen Tangle, sondern aus der Tangle-Struktur des Graphen: Wir zeigen, dass sich das Tree-of-Tangles Theorem von Robertson und Seymour auf lokal endliche Graphen ohne dicke Enden erweitern lässt, indem wir an jeder Obstruktion für die Existenz solcher Tangle-unterscheidender Baumzerlegungen dicke Enden (und damit Halbgitter-Minoren) finden.

In Teil III untersuchen wir, wie Baumzerlegungen mit verschiedenen Konzepten von Zusammenhang in Graphen wechselwirken. Insbesondere konstruieren wir eine Baumzerlegung, die die verschiedenen Typen von Unendlichkeiten in unendlichen Graphen darstellt, und erweitern den Satz von Menger auf bigerichtete Graphen. Ein wiederkehrendes Thema sind dabei strukturelle Dualitäten im Hinblick auf Weitemaße und deren jeweilige Obstruktionen, wie beispielsweise das Tangle-Tree Duality Theorem und die verschiedenen Gittersätze.

In Kapitel 7 beweisen wir, dass jeder Graph, der eine Baumzerlegung in endliche Teile hat, auch eine gewurzelte Baumzerlegung in endliche Teile besitzt, die verbunden, eng und komponententreu ist. In Kapitel 8 argumentieren wir, dass dieses Resultat das Beste ist, was man sich erhoffen kann, indem wir Gegenbeispiele für diverse allgemeinere oder stärkere Aussagen konstruieren.

Als erste Anwendung des Hauptresultats von Kapitel 7 erhalten wir, dass jeder Graph ohne Halbgitter-Minor eine schlanke Baumzerlegung in endliche Teile besitzt; dies verstärkt das entsprechende Ergebnis von Kříž und Thomas für Graphen mit endlich beschränkter Baumweite. Insbesondere folgt daraus, dass jeder Graph ohne Halbgitter-Minor eine Baumzerlegung besitzt, die all seine Enden und kritischen Eckenmengen effizient unterscheidet; dies verstärkt die Ergebnisse von Carmesin sowie von Elm und Kurkofka für diese Graphklasse.

Als zweite Anwendung zeigen wir, dass jeder Graph G, der eine Baumzerlegung in endliche Teile hat, auch eine Baumzerlegung in endliche Teile besitzt, die alle Unendlichkeiten von G darstellt, d. h. sie stellt die Enden von G homöomorph dar mit ihren kombinierten Graden und ihren dominierenden Ecken (womit wir eine Frage von Halin aus dem Jahre 1977 beantworten) sowie die kritischen Eckenmengen von G und ihre engen Komponenten kofinit. Letztere Baumzerlegung ermöglicht einheitliche Beweise für die Charakterisierungen von Graphen ohne Halbgitter-Minor und von Graphen ohne Unterteilung eines unendlichen Binärbaums, die auf Robertson, Seymour und Thomas zurückgehen. Darüber hinaus verwendete Albrechtsen diese Baumzerlegung, um das Tangle-Tree Duality Theorem von Robertson und Seymour auf unendliche Graphen zu verallgemeinern.

In Kapitel 9 beweisen wir eine exakte strukturelle Dualität für starke Immersionsminoren von Mauern anhand spezieller Baumschnittzerlegungen. Für den Fall von Mauern verbessert dies den Struktursatz von Dvořák und Wollan für Graphen ohne einen festen Graphen als Immersionsminor; dieser Satz hat für Mauern keine Umkehrung.

In Kapitel 10 erweitern wir den Satz von Lovász-Cherkassky von lokal endlichen Graphen auf lokal endliche Graphen mit ihren Enden. Für eine Menge T von Ecken eines endlichen Graphen betrachtet der Satz von Lovász-Cherkassky kantendisunkte T-Pfade und zeigt,

dass eine natürliche obere Schranke angenommen wird, wenn alle Ecken außerhalb von T geraden Grad haben. Jóo verallgemeinerte dieses Resultat auf lokal endliche Graphen, und wir verallgemeinern es weiter auf den Fall, in dem T sowohl Ecken als auch Enden enthalten kann, wobei Pfade unendlich werden, wenn sie ein Ende verbinden.

In Kapitel 11 untersuchen wir, inwiefern sich der Satz von Menger auf bigerichtete Graphen (eine Verallgemeinerung von (un)gerichteten Graphen, bei der jede Kante an jedem ihrer Endecken das Vorzeichen + oder − hat) übertragen lässt. Wir zeigen, dass der Satz von Menger in bigerichteten Graphen im Allgemeinen nicht gilt. Nichtsdestotrotz finden wir eine hinreichende Bedingung für X und Y, die eine Version des Satzes von Menger in bigerichteten Graphen liefert, die insbesondere dessen (un)gerichtete Pendants verallgemeinert.

In Kapitel 12 beweisen wir, dass für jede Menge S von Ecken eines gerichteten Graphen D die maximale Anzahl von Ecken in S, die in einer Familie von eckendisjunkten Kreisen in D enthalten sind, mindestens die minimale Größe einer Eckenmenge ist, die alle Kreise trifft, die eine Ecke von S enthalten. Wir zeigen außerdem, dass die entsprechende Aussage auch in bigerichteten Graphen wahr ist und dass ihre Kantenvariante sowohl in ungerichteten als auch in gerichteten Graphen gilt, aber nicht in bigerichteten Graphen; die Eckenvariante in ungerichteten Graphen bleibt eine offene Frage. Als Anwendung leiten wir daraus ab, dass die (bi)gerichtete Baumweite eines (bi)gerichteten Graphen linear durch seine Kreisweite begrenzt ist, was die bisher bekannte quadratische Schranke verbessert.

In Teil IV betrachten wir die grobe Geometrie von Graphen mithilfe von Werkzeugen und Methoden der Zusammenhangstheorie. Der Begriff „grobe Graphentheorie“ wurde von Georgakopoulos und Papasoglu eingeführt, die Gromovs Ansatz der groben Geometrie auf Graphen übertrugen und eine Sammlung von Problemen und Resultaten zum Zusammenspiel von Graphenminoren und Geometrie präsentierten. Wie bei Gromov werden auch bei Georgakopoulos und Papasoglu Quasiisometrien zum zentralen Werkzeug: Sie bieten das formale Framework für die Untersuchung der groben Geometrie von Graphen.

In Kapitel 13 and 14 beweisen wir mehrere Fälle der folgenden Vermutung von Georgakopoulos und Papasoglu über „fette Minoren“, eine metrische Variante von Minoren: *Sei $\mathcal{X}$ eine endliche Menge endlicher Graphen. Dann gibt es eine Funktion $f : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$, sodass für alle $K \in \mathbb{N}$ und jeden Graphen G die folgende Aussage gilt: Wenn G kein $X \in \mathcal{X}$ als K-fetten Minor enthält, dann ist G $f(K)$-quasi-isometrisch zu einem Graphen, der kein $X \in \mathcal{X}$ als Minor enthält.* Unsere Resultate basieren auf einer neuen Perspektive auf Quasiisometrien mithilfe von Graphenzerlegungen: Wir zeigen, dass Quasiisometrien mit beschränkten Parametern äquivalent durch ehrliche Graphenzerlegungen mit beschränkter radialer Weite und radialer Ausdehnung beschrieben werden können.

In Kapitel 13 beweisen wir die obige Vermutung für $\mathcal{X} = \{K^3, K_{1,3}\}$ (quasi-isometrisch zu einer disjunkten Vereinigung von Pfaden), $\mathcal{X} = \{K_{1,3}\}$ (Kreise und Pfade) und $\mathcal{X} = \{K^3, W\}$ (unterteilte Sterne). In all diesen Fällen gelten unsere Resultate sogar dann, wenn man „fette Minoren“ durch den stärkeren Begriff der „quasi-geodätischen topologischen Minoren“ ersetzt, die man sich als metrisches Analogon topologischer Minoren vorstellen kann.

In Kapitel 14 bestätigen wir die „fette Minoren“-Vermutung für $\mathcal{X} = \{K^4\}$. Unsere

Beweistechnik liefert zudem einen neuen kurzen Beweis für den entsprechenden K_4^--Fall, der zuvor von Fujiwara und Papasoglu bewiesen worden war.

In Kapitel 15 beweisen wir eine metrische Version des Satzes von Menger für zwei Wege. Wir zeigen, dass für alle $d \in \mathbb{N}$ und für alle Graphen G und Eckenmengen $X, Y \subseteq V(G)$ folgende Aussage gilt: Entweder gibt es 2 disjunkte X–Y-Pfade in G, die paarweise Abstand mindestens d haben, oder es gibt eine Ecke $v \in V(G)$, sodass jeder X–Y-Pfad Abstand höchstens $129d$ zu v hat. Dieses Ergebnis dient auch als wichtiges Tool in Kapitel 14.

APPENDIX C

Publications related to this thesis

Part I

Chapter 1 is based on [DJKK25; DJKK26].
Chapter 2 is based on [JKK25].
Chapter 3 is based on [JK25].

Part II

Chapter 4 is based on [ABJKW24].
Chapter 5 is based on [DEJ24].
Chapter 6 is based on [JK24].

Part III

Chapter 7 is based on [AJKP24].
Chapter 8 is based on [AJKP25].
Chapter 9 is based on [DJKW25].
Chapter 10 is based on [JJKKM23].
Chapter 11 is based on [BGGJR23].
Chapter 12 is based on [BGGJR24].

Part IV

Chapter 13 is based on [ADEFJKW23].
Chapter 14 is based on [AJKW25].
Chapter 15 is based on [AHJKW24].

APPENDIX D

Declaration of contributions

The research in this thesis is based on work that I have done with various co-authors (see also Appendix 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. In cases where I do not emphasise who did which work, it should be assumed that all involved authors contributed equally to the research.

Part I

Chapter 1 is based on [DJKK26], which is joint work with Reinhard Diestel, Paul Knappe and Jan Kurkofka. The basic idea of using canonical tree-decompositions of local coverings is due to Reinhard Diestel and Johannes Carmesin. The whole theory in all its details was then developed by Paul Knappe, Jan Kurkofka and me. In particular, Paul Knappe and I 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 [JKK25], which is joint work with Paul Knappe and Jan Kurkofka. The overarching idea towards a theory of local separations was developed by Paul Knappe, Jan Kurkofka and me. Displacement as a sufficient criterion for settings in which the theory of local coverings works particularly well is due to Paul Knappe and me as well as the respective classes of examples Section 2.5. Paul Knappe, Jan Kurkofka and I 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 Paul Knappe, Jan Kurkofka and me.

Chapter 3 is based on [JK25], which is joint work with Paul Knappe. Paul Knappe and I developed most results in this chapter together; the proof of Theorem 9 is by Paul Knappe. A first presentation was worked out by Paul Knappe, while the final presentation is due to me.

Part II

Chapter 4 is based on [ABJKW24], which is joint work with Sandra Albrechtsen, Hanno von Bergen, Paul Knappe and Paul Wollan. We have in close collaboration developed a rough outline of the proof during a 3-weeks research visit of Paul Wollan in Hamburg in October 2021. Sandra Albrechtsen and Hanno von Bergen developed a more detailed outline of the proof. Afterwards, Sandra Albrechtsen, Paul Knappe and I worked out the full details of the proofs and developed the final presentation.

Chapter 5 is based on [DEJ24], which is joint work with Reinhard Diestel and Christian Elbracht. Reinhard Diestel started this line of research and also came up with the notion of the resilience as a potential sufficient criterion for the existence of inducing point sets. Christian Elbracht then proved that resilience indeed forms such a sufficient criterion (Theorem 5.3.1), which I later extended to the criterion of being locally induced (Theorem 5.3.5). I proved the results on orientations and duality of set separations Section 5.4. Christian Elbracht showed that extendable tangles admit inducing point sets Section 5.5. The write-up of all results is mostly due to me.

Chapter 6 is based on [JK24], which is joint work with Paul Knappe. Paul Knappe and I developed all the proofs and their details together as well as the complete presentation.

Part III

Chapter 7 is based on [AJKP24], which is joint work with Sandra Albrechtsen, Paul Knappe and Max Pitz. Max Pitz started the project by asking Paul Knappe and me a question about efficiently distinguishing ends in a graph. Based on our 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. Paul Knappe generalised the algorithm to produce, for any 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 (including Sandra Albrechtsen), 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 parts to obtain the final tree-decomposition into finite parts (Section 7.4.3). Theorem 20 and Corollary 22 are due to Sandra Albrechtsen and Corollaries 7.1.2 and 7.1.3 are due to Max Pitz. Sandra Albrechtsen, Paul Knappe and I worked out the details of all proofs in this chapter. The final presentation, which included a complete overhaul of the story, was then developed by Sandra Albrechtsen and Paul Knappe.

Chapter 8 is based on [AJKP25], which is joint work with Sandra Albrechtsen, Paul Knappe and Max Pitz. Sandra Albrechtsen developed Construction 8.3.1 and proved Examples 25 and 26. For Example 27, Max Pitz suggested considering the graph from [C19, Example 3.7], and Sandra Albrechtsen proved that it indeed works. Paul Knappe,

Max Pitz and I developed Theorem 8.5.1. Paul Knappe and Sandra Albrechtsen worked out the details of the proofs.

Chapter 9 is based on [DJKW25], which is joint work with Reinhard Diestel, Paul Knappe 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. Paul Knappe and I then worked out all details and developed a complete presentation.

Chapter 10 is based on [JJKKM23], which is joint work with Attila Joó, Paul Knappe, Jan Kurkofka and Ruben Melcher. The authors developed the proof idea during a series of online meetings. Paul Knappe, Jan Kurkofka and I then worked out the details and developed the complete presentation.

Chapter 11 is based on [BGGJR23], which is joint work with Nathan Bowler, Ebrahim Ghorbani, Florian Gut and Florian Reich. We conducted all research together in regular meetings. The first detailed drafts were written by Nathan Bowler Sections 11.3 and 11.5, Florian Reich Sections 11.4 and 11.7 and Florian Gut Section 11.6. I revised all drafts and developed the final presentation together with Florian Reich and Florian Gut.

Chapter 12 is based on [BGGJR24], which is joint work with Nathan Bowler, Ebrahim Ghorbani, Florian Gut and Florian Reich. We came up with this question together, aiming for a simpler proof of Corollary 12.1.5. Theorem 12.2.1 was proved by Ebrahim Ghorbani. Based on a detailed proof idea of Nathan Bowler and Irene Muzi, we showed Theorem 12.2.2 together. We jointly developed proof ideas towards Theorem 12.2.7 and Theorem 12.2.3, and Ebrahim Ghorbani finally made Theorem 12.2.3 work. The counterexample Proposition 12.2.8 is due to Florian Reich. The applications including Corollary 12.1.5 are joint work. The final presentation of the chapter was mostly developed by Florian Reich and me.

Part IV

Chapter 13 is based on [ADEFJKW23], which is joint work with Sandra Albrechtsen, Reinhard Diestel, Ann-Kathrin Elm, Eva Fluck, Paul Knappe and Paul Wollan. During a 2-weeks research visit of Paul Wollan and Eva Fluck in Hamburg in April 2022, we developed in close collaboration a detailed outline of a proof of a version of Theorems 33 and 34, where we excluded 3-geodesic subdivision of K_3. Ann-Kathrin Elm came up with and proved Lemma 13.4.4, which improved our earlier version of Theorem 33 to the one presented in this thesis. Sandra Albrechtsen proved Theorem 35. Sandra Albrechtsen, Ann-Kathrin Elm, Eva Fluck, Paul Knappe and I worked out the details of the proofs and finalised a first version (arXiv:2307.08497v1). The final presentation of the results in Sections 13.1 and 13.3 in this thesis, which takes into account several recent developments in coarse graph theory, is due to Sandra Albrechtsen, Paul Knappe and me.

Chapter 14 is based on [AJKW25], which is joint work with Sandra Albrechtsen, Paul Knappe and Paul Wollan. During a 3-weeks research visit of Paul Wollan in Hamburg in June and July 2022, we gathered in close collaboration various proof ideas for Theorems 37 and 38, but were unsuccessful in our attempts. Sandra Albrechtsen then proved Theorems 37

and 38 building on some of these ideas; in particular, the proof of Lemma 14.5.2 is completely by Sandra Albrechtsen. Paul Knappe formalised the inductive proof method described in Section 14.3. Sandra Albrechtsen and Paul Knappe worked out the details of the proofs of Theorems 37 and 38 and developed the final presentation.

Chapter 15 is based on [AHJKW24], which is joint work with Sandra Albrechtsen, Tony Huynh, Paul Knappe and Paul Wollan. This project was started during a 3-weeks research visit of Paul Wollan in Hamburg in June and July 2022. Sandra Albrechtsen, Paul Knappe, Paul Wollan and I have in close collaboration developed Conjecture 41 and tried various proof ideas for Theorem 40, but were unsuccessful in our attempts. Tony Huynh and Paul Wollan then proved Theorem 40 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 441. Sandra Albrechtsen, Paul Knappe and I worked out the details of this fix and rewrote the construction of the two paths P_1, P_2 starting on page 441.

APPENDIX E

Acknowledgements

I am deeply grateful to everyone who supported me along the way. This dissertation would not have been possible without so many different people.

I thank my supervisor, Reinhard Diestel, for making this dissertation possible and for his support, both academically and personally, throughout all the ups and downs. I am especially grateful for his guidance in identifying research topics and questions worth pursuing and for his continuous advice on writing.

I am grateful to all my co-authors, especially my office mates Paul Knappe and Sandra Albrechtsen as well as my close collaborator Jan Kurkofka, for the many adventures we shared, for your open ears, your patience, and the many enjoyable hours together. I also thank all my colleagues in the Discrete Mathematics research group in Hamburg for the pleasant, open working atmosphere and the spirit of team work.

My thanks also go to the Studienstiftung and the Cusanuswerk for supporting me throughout my studies and doctoral work. Their support not only facilitated this path financially, but the many encounters and opportunities they enabled helped me grow personally.

Last but not least, I owe my deepest thanks to my partner Annika Hohm, to my family, Christiane and Daniela Jacobs as well as Immanuel Jacobs (in loving memory), and my friends and friend family – you know who you are :) Without your love and support over so many years, much would not have been possible, including this dissertation, and life would be only half as joyful.

APPENDIX F

Eidesstattliche Versicherung

Eidesstattliche Versicherung
Hiermit versichere ich an Eides statt, die vorliegende Dissertationsschrift selbst verfasst und keine anderen als die angegebenen Hilfsmittel und Quellen benutzt zu haben. Sofern im Zuge der Erstellung der vorliegenden Dissertationsschrift generative Künstliche Intelligenz (gKI) basierte elektronische Hilfsmittel verwendet wurden, versichere ich, dass meine eigene Leistung im Vordergrund stand und dass eine vollständige Dokumentation aller verwendeten Hilfsmittel gemäß der Guten wissenschaftlichen Praxis vorliegt. Ich trage die Verantwortung für eventuell durch die gKI generierte fehlerhafte oder verzerrte Inhalte, fehlerhafte Referenzen, Verstöße gegen das Datenschutz- und Urheberrecht oder Plagiate.

Affidavit
I hereby declare and affirm that this doctoral dissertation is my own work and that I have not used any aids and sources other than those indicated. If electronic resources based on generative artificial intelligence (gAI) were used in the course of writing this dissertation, I confirm that my own work was the main and value-adding contribution and that complete documentation of all resources used is available in accordance with good scientific practice. I am responsible for any erroneous or distorted content, incorrect references, violations of data protection and copyright law or plagiarism that may have been generated by the gAI.

Hamburg, den 09.05.2026

Raphael W. Jacobs