

Connectivity, tree-decompositions and unavoidable-minors

by

Benson Joeris

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Combinatorics and Optimization

Waterloo, Ontario, Canada, 2015

© Benson Joeris 2015

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The results in this thesis are steps toward bridging the gap between the handful of exact structure theorems known for minor-closed classes of graphs, and the very general, yet wildly qualitative, Graph Minors Structure Theorem.

This thesis introduces a refinement of the notion of tree-width. Tree-width is a measure of how “tree-like” a graph is. Essentially, a graph is tree-like if it can be decomposed across a collection of non-crossing vertex-separations into small pieces. In our variant, which we call *k-tree-width*, we require that the vertex-separations each have order at most k .

Tree-width and branch-width are related parameters in a graph, and we introduce a branch-width-like variant for k -tree-width. We find a dual notion, in terms of tangles, for our branch-width parameter, and we prove a generalization of Robertson and Seymour’s Grid Theorem.

Acknowledgements

First, I would like to thank my supervisor Jim Geelen for his patience and dedication over the last several years. He has been a constant source of interesting open problems and insightful ideas.

I would also like to thank my committee members—Bruce Richter, Bertrand Guenin, Anna Lubiw and Guoli Ding—for their careful reading of this thesis and very helpful comments.

I would like to thank my undergraduate advisor Ross McConnell for seeing the potential in this bored high school and helping me to discover the beautiful world of mathematics research.

I would like to thank my friends, particularly Victor and Yomna, for their support, and for the joy we have shared playing countless board games.

I would like to thank my family for their love and for encouraging me to find my path.

Finally, I would like to thank my partner Isabel. Without her patient, loving support through the ups and downs over the years, I would not have completed this thesis.

Dedication

To Isabel, the best result of my Ph.D.

Table of Contents

1	Introduction	1
1.1	Exact structure theorems	1
1.2	Graph Minors Structure Theorem	4
1.3	Separations and connectivity	7
1.4	Graphs with no K_6 -minor	10
1.4.1	Apex graphs	10
1.4.2	Planar plus a triangle	11
1.4.3	Doublecross	12
1.4.4	Hose	12
1.4.5	Hamburger	13
1.4.6	A conjecture	15
1.5	Tree-decompositions	16
1.5.1	Clique-sums	18
1.5.2	Tree-width and branch-width	22
1.5.3	θ -tree-width and θ -branch-degree	23
1.5.4	Displaying a minor	24
1.6	Tangles	26
1.6.1	Duality theorems	28
1.6.2	Tree of tangles	29

1.7	θ -connected sets	30
1.8	Minors arising from θ -tangles	31
1.8.1	The main theorem	32
1.8.2	Specialization for $\theta = 6$	34
1.9	Summary of results	34
2	Tree-decompositions and tangles	37
2.1	Introduction	37
2.2	Connectivity systems	38
2.2.1	Examples of connectivity systems	39
2.2.2	Domination	39
2.3	Tree-decompositions	40
2.4	Tangles	41
2.4.1	Induced tangles	42
2.5	Covering-number	43
2.6	The duality theorem	45
2.7	Tangle matroids	49
2.8	Graphic tangles and their matroids	52
2.9	θ -connected sets	52
2.9.1	Uniform matroids	54
2.9.2	Hyperplanes in matroids	54
2.9.3	Constructing a θ -connected set	55
3	Trees of tangles	57
3.1	Introduction	57
3.2	Displaying tangles	58
3.3	Distinguishing separations	59
3.4	Tree of tangles	60

3.4.1	Bounded-degree tree of θ -tangles	62
3.5	Partial tree-decomposition	63
3.6	Partial tree of θ -tangles	64
3.7	Robust sets	65
3.8	Tiebreakers	69
3.9	Distinguishing covering	70
3.10	Constructing a bounded-degree tree of θ -tangles	72
4	Unavoidable-minors	78
4.1	Introduction	78
4.1.1	Unavoidable-minors	78
4.1.2	θ -connected sets	82
4.1.3	The main theorem	83
4.2	Homogeneous wheels	84
4.3	Connectivity and tangles in wheels	88
4.3.1	Weak θ -connectivity	92
4.3.2	Large θ -connected sets	93
4.4	Alignment of minors and tangles	94
4.5	Necklace	96
4.5.1	Basic properties of beads and necklaces	99
4.6	Initial necklace	102
4.6.1	From a tangle to disjoint connected subgraphs	102
4.6.2	From disjoint connected subgraphs to an initial necklace	106
4.7	Jumps in a necklace	111
4.7.1	Jump-free necklaces	113
4.7.2	Necklaces with jumps to the weak bead	119
4.8	Growing a necklace	127
4.9	Necklace to wheel	129

4.10 Grid Theorem	135
4.10.1 Constructing homogeneous wheels	135
4.10.2 Tangle truncation	136
4.10.3 Proof of Grid Theorem	137
5 Conclusions and Further Work	140
Bibliography	144
Notation	147
Index	148

Chapter 1

Introduction

The results in this thesis are motivated by the following question:

Question. What is the structure of graphs with no K_6 -minor?

The structure of graphs with no K_5 -minor is described exactly by a classic theorem of Wagner (Theorem 1.1.3), but the structure of graphs with no K_6 -minor is known only very roughly. We do not make direct progress on an exact description of the structure of graphs with no K_6 -minor, but we introduce a new qualitative approach that improves substantially on existing techniques.

The results presented here are quite general and may also be applicable to related problems, such as understanding the structure of graphs with no Petersen Graph-minor, but we will focus our discussion on K_6 .

Our main result, Theorem 1.8.2, is a generalization of Robertson and Seymour’s Grid Theorem [23]; it characterizes graphs containing large, “ θ -connected pieces” by a class of highly structured minors that such graphs must contain.

1.1 Exact structure theorems

A *minor* of a graph G is a graph G' obtained from G by a series of vertex-deletions, edge-deletions and edge-contractions—an *edge-contraction* of an edge e incident with vertices u and v in graph G is a graph G/e obtained from G by identifying vertices u and v and deleting the edge e . If G and H are graphs, such that G has a minor isomorphic to H , then




Figure 1.1: G is a 3-sum between graphs G_1 and G_2 along the 3-vertex cliques H_1 and H_2 .

we say that G has an H -minor, while if G has no minor isomorphic to H , then we say that G is H -minor-free. We are interested in graphs that are K_6 -minor-free, but let us begin with some easier related questions about graphs that do not have other small minors.

A graph with no K_3 -minor cannot have any cycles of length greater than 2, and hence must be a forest (possibly with some parallel edges and/or loop edges). This means that graphs with no K_3 -minor can be built up from graphs on at most 2 vertices by two simple operations—a 0-sum, which is just the disjoint union of two graphs, and a 1-sum which takes two vertex-disjoint graphs and identifies exactly one vertex from the first graph with exactly one vertex from the second graph.

Proposition 1.1.1. *A graph has no K_3 -minor if and only if it can be constructed from graphs with at most two vertices by 0-sums and 1-sums.*

For graphs that exclude cliques larger than K_3 , we must generalize the 0-sum and 1-sum operations as follows. A k -sum, or a *clique-sum of order k* , combines two vertex-disjoint graphs G_1 and G_2 with complete subgraphs H_1 and H_2 in G_1 and G_2 respectively, where $|V(H_1)| = |V(H_2)| = k$, by bijectively identifying each vertex in H_1 with a vertex in H_2 , and removing all of the edges in $E(H_1) \cup E(H_2)$. See Figure 1.1.

Graphs with no K_4 -minor are called *series-parallel graphs* [8], and can be constructed using clique-sums of order at most 2.




Figure 1.2: The graph V_8

Proposition 1.1.2. *A graph has no K_4 -minor if and only if it can be constructed from graphs on at most three vertices by 0-sums, 1-sums and 2-sums.*

A classic theorem of Wagner [31] (see also [30]) shows that graphs with no K_5 -minor possess a similar structure. The graph V_8 is obtained from an eight vertex cycle by adding the 4 diagonal edges; see Figure 1.2.

Theorem 1.1.3 (Wagner’s Theorem). *A graph has no K_5 -minor if and only if it can be obtained from planar graphs and copies of V_8 by clique-sums of order at most three.*

Proposition 1.1.1, Proposition 1.1.2 and Theorem 1.1.3 are all *exact structure theorems*—the sets of graphs built using the respective constructions are exactly the sets of graphs excluding the respective minors. Exact structure theorems are also known for a few other small graphs.

Kelmans [18] and, independently, Robertson [22], described the structure of graphs with no V_8 -minor.

Theorem 1.1.4. *A graph G has no V_8 -minor if and only if G can be obtained by clique-sums of order at most three from graphs in a family \mathcal{H} such that, for each $H \in \mathcal{H}$, either*

1. H is planar,
2. H has two distinct vertices x and y such that $H - \{x, y\}$ is a cycle,
3. H has a set B of four vertices such that every edge in H is incident with a vertex in B ,




Figure 1.3: The cube graph (a) and the octahedron graph (b)

4. H is isomorphic to the line graph of $K_{3,3}$, or
5. $|V(H)| \leq 7$.

Maharry [19] gave an exact structure theorem for graphs with no cube-minor; the cube graph is pictured in Figure 1.3a. Ding [5] recently described the structure of graphs with no octahedron-minor; the octahedron graph is pictured in Figure 1.3b. Ding and Liu [7, 6] recently gave an exact structure theorem of graphs with no H -minor for each *3-connected* graph H —a graph H is *3-connected* if the deletion of any two vertices in H leaves a connected graph) with 11 edges; together with previous results, exact structure theorems for graphs with no H -minor are known for all 3-connected graphs with at most 11 edges.

The only 3-connected graphs with more than 11 edges for which exact structure theorems are known are V_8 , the cube and the octahedron, which each have 12 edges. On the other hand, K_6 has 15 edges, so it does not seem likely that existing techniques will yield an exact structure theorem for graphs with no K_6 -minor any time soon.

1.2 Graph Minors Structure Theorem

With an exact structure theorem out of reach, we turn our attention to *qualitative* methods—that is, we seek a graph construction such that each graph with no K_6 -minor can be constructed, and each graph that can be constructed is “close to” a graph with no K_6 -minor. This can be seen as a kind of “approximation” of the structure of graphs with

no K_6 -minor, allowing us to prove useful structural results despite the apparent difficulty in finding an exact description. Of course, we want our notion of a graph being “close to” a graph with no K_6 -minor to be as tight as possible, to approximate the exact structure as closely as possible.

Robertson and Seymour’s Graph Minors Structure Theorem [26] provides, for each graph H , a very rough qualitative description of the structure of graphs with no H -minor in terms of graph *embeddings*—a graph *embeds* in a surface if it can be drawn on the surface with no edge crossings. The precise statement of the Graph Minors Structure Theorem is technical and not necessary for our purposes; it essentially says that, for each graph H , a graph with no H -minor can be constructed using clique-sums from graphs which “almost” embed in some surface in which H cannot be embedded, where the measure of “almost” embedding in the surface is bounded by a constant depending only on the graph H .

For the special case of $H = K_6$, the Graph Minors Structure Theorem becomes much simpler. The only surface on which K_6 cannot be embedded is the plane, and notion of “almost” embedding in the plane can be simplified to the notion of *k-apex*—a *k-apex* graph is a graph G with a set X of at most k vertices such that $G - X$ is planar.

Theorem 1.2.1. *There exists a constant $N \in \mathbb{N}$ such that if G is a graph with no K_6 -minor, then G can be constructed by clique-sums from N -apex graphs.*

This description of the structure of K_6 -minor-free-graphs is extremely coarse compared to the exact results in Section 1.1. The constant N in Theorem 1.2.1 is astronomical, so, for example, K_{1000} is one of the pieces that can be used to construct a graph with no K_6 -minor.

What is even more troubling is that these large cliques can then be used in clique-sum operations, so the order of the clique-sums used in the construction can be huge. A K_6 -minor can have its edges distributed in a totally arbitrary way between the two sides of a clique sum of order 6 or greater, as shown in Figure 1.4a. This means that the graphs constructed by clique sums from N -apex graphs can have K_6 -minors that do come from any particular graph used in the construction, but emerge only in the global structure. This is in contrast with constructions using clique-sums of order at most 5, where any particular K_6 -minor must have at least one vertex that is incident only with edges from one side of the clique-sum, as shown in Figure 1.4b. This means that a K_6 -minor in a graph constructed by clique-sums of order at most 5 can be associated in a natural way with a K_6 -minor in exactly one of the graphs used in the construction.

Also, K_5 and K_6 are indistinguishable by this kind of structure theorem—for K_5 -minor-free graphs, the Graph Minors Structure theorem gives a result of an identical form to Theorem 1.2.1, but with a different constant N .




Figure 1.4: A K_6 -minor in a graph constructed by a clique-sum of order 6 (a) and order 5 (b), where contracting the dashed edges between vertices with the same label yields K_6 . The K_6 -minor can be distributed arbitrarily between the two sides of a clique-sum of order 6 or greater (a), but must lie primarily on one side of a clique-sum of order at most 5, which is the side that all of the edges incident with one vertex of the K_6 (b).

The results in this thesis are designed to help prove a more refined qualitative structure theorem for graphs with no K_6 -minor.

1.3 Separations and connectivity

A *separation* in a graph G is a bipartition of the edges of G ; the *order* of a separation (A, B) is the number of vertices in G incident with both an edge in A and an edge in B ; for $\theta \in \mathbb{N}$, a separation of order at most θ is called a θ -*separation*. For $\theta \in \mathbb{N}$, a graph G is θ -*connected* if, for every $(\theta - 1)$ -separation (A, B) , of G , either every vertex in G is incident with an edge in A , or every vertex in G is incident with an edge in B . Note that we do *not* require that a θ -connected graph have more than θ vertices.

For small values of θ , θ -connectivity is related to clique-sum constructions of order at most $(\theta - 1)$ as follows. If G is a graph that is not (1) -connected, then G can be expressed as a 0-sum of two proper subgraphs of G ; recursively, this shows that each graph G can be constructed using 0-sums from connected subgraphs of G . If G is a connected graph that is not 2-connected, then G can be expressed as a 1-sum of two connected, proper subgraphs of G by splitting a cut-vertex; therefore, each graph G can be constructed using 0-sums and 1-sums from 2-connected subgraphs of G . If G is a 2-connected graph that is not 3-connected, then G can be expressed as a 2-sum of two 2-connected, proper *minors* of G —in a 2-connected graph, each side of a separation (A, B) of 2 is connected, and hence can be contracted to a single edge between the two vertices, each incident with both an edge in A and an edge in B ; see Figure 1.5; this implies the following lemma.

Lemma 1.3.1. *Each graph G can be constructed from 0-sums, 1-sums and 2-sums from 3-connected minors of G .*

Thus, any graph G with no K_4 -minor can be constructed using 0-sums, 1-sums and 2-sums from 3-connected minors of G , each of which has no K_4 -minor, so the earlier exact structure theorem for graphs with no K_4 -minors, Proposition 1.1.2, is equivalent to the following.

Proposition 1.3.2. *If G is a 3-connected graph with no K_4 -minor, then $|V(G)| < 4$.*

One might hope that this pattern continues—that graphs can be constructed from their θ -connected minors by clique-sums of order at most $\theta - 1$ —but this begins to break down for $\theta > 3$. Suppose G is a graph with a vertex u with three neighbours, v_1 , v_2 and v_3 along edges e_1 , e_2 and e_3 respectively, and G has no edges between vertices in $\{v_1, v_2, v_3\}$. Then




Figure 1.5: An illustration of the fact that a 2-connected graph with a separation of order 2 is a 2-sum of two minors; the bold path on each side of the separation can be contracted to an edge between the two cut vertices.

$(\{e_1, e_2, e_3\}, E(G) - \{e_1, e_2, e_3\})$ is a 3-separation, demonstrating that G is not 3-connected. The corresponding 3-sum construction of G , however, involves a graph isomorphic to K_4 and a graph, G' , constructed from G by deleting u and adding edges between each pair of vertices in $\{v_1, v_2, v_3\}$, as shown in Figure 1.6. The graphs G and G' have the same number of edges, so G' is not a proper minor of G .

It turns out that this problem with degree 3 vertices is the only thing that can cause a graph to have a 3-separation but not a 3-sum construction, so the situation can be partially salvaged by relaxing our notion of connectivity. A graph G is *internally 4-connected* if it is 3-connected and, for each separation (A, B) of order 3, $\min\{|A|, |B|\} = 3$.

Lemma 1.3.3. *Each graph G can be constructed using clique-sums of order at most 3 from internally 4-connected minors of G .*

Thus, any graph G with no K_5 -minor can be constructed using clique-sums of order at most 3 from internally 4-connected minors of G , each of which has no K_5 -minor, so Wagner's structure theorem for graphs with no K_5 -minor, Theorem 1.1.3, is equivalent to the following.

Theorem 1.3.4 (Wagner's Theorem, connectivity version). *An internally 4-connected graph G has no K_5 -minor if and only if G is either planar or isomorphic to V_8 .*




Figure 1.6: G is a graph with a 3-separation, but the clique-sum corresponding to this 3-separation involves a graph G' that is *not* a minor of G .

Because internal 4-connectivity is weaker than 4-connectivity, and V_8 is internally 4-connected but not 4-connected, Theorem 1.3.4 implies the following relationship between K_5 and planar graphs.

Theorem 1.3.5. *Every 4-connected graph with no K_5 -minor is planar.*

Jørgensen conjectured [14] that a similar relationship holds between K_6 and *apex graphs*—that is, graphs which possess some vertex whose deletion yields a planar graph.

Conjecture 1.3.6 (Jørgensen’s Conjecture). *Every 6-connected graph with no K_6 -minor is apex.*

Kawarabayashi, Norine, Thomas and Wollan recently proved [17] that Jørgensen’s Conjecture is true for sufficiently large 6-connected graphs.

Theorem 1.3.7. *There exists a constant $N \in \mathbb{N}$ such that every 6-connected graph on at least N vertices with no K_6 -minor is apex.*

Norine and Thomas have also announced [17] that Theorem 1.3.7 can be generalized as follows, although the proof has not yet been published.

Theorem 1.3.8. *For each integer t , there exists an integer N_t such that every t -connected graph on at least N_t vertices with no K_t -minor is $(t - 5)$ -apex.*

Whether or not Jørgensen’s Conjecture holds for all 6-connected graphs, Theorem 1.3.7 suggests that the relationship between K_6 and apex graphs may be similar to the relationship between K_5 and planar graphs.

This analogy, while quite promising, does have some important limitations. First, the constant N in Theorem 1.3.7 is quite large, so, if we use this to describe the structure of graphs with no K_6 minors as being constructed from apex graphs and a finite family of “small graphs”, as in Wagner’s Theorem, then the “small graphs” will be much larger than V_8 . This limitation seems hard to avoid, and, as we are seeking a *qualitative* as opposed to *exact* structure theorem, we are prepared to accept it.

The larger issue, however, is that the structural version of Wagner’s Theorem, Theorem 1.1.3, is equivalent to the connectivity version, Theorem 1.3.4. As we saw, decomposing along 3-sums into 4-connected pieces is already problematic, requiring the use of *internal* 4-connectivity instead, and decomposing along 5-sums into “6-connected pieces”, is even more problematic. Therefore, Theorem 1.3.7 does not by itself lead to a structure theorem for graphs with no K_6 -minor. In this thesis we describe a way that graphs can be decomposed into “weakly 6-connected pieces” (something weaker than a 6-connected graph). Because we need to relax the notion of connectivity, Theorem 1.3.7 cannot be immediately applied, but the overall strategy seems promising.

1.4 Graphs with no K_6 -minor

Here we review some natural classes of graphs that do not contain K_6 -minors.

1.4.1 Apex graphs

Conjecture 1.3.6 and Theorem 1.3.7 suggest that apex graphs, the class of graphs which can be constructed by adding a single vertex to a planar graph, as in Figure 1.7, are a fundamental class of K_6 -minor-free graphs.

Proposition 1.4.1. *If a graph is apex, then it does not contain a K_6 -minor.*

Proof. The class of apex graphs is closed under taking minors, and K_6 is not apex. □




Figure 1.7: A schematic of the construction of an apex graph; deleting vertex v_0 leaves a graph embedded in the plane, represented by the shaded region.




Figure 1.8: A schematic of the construction of planar-plus-triangle graphs; the shaded region represents a graph embedded in the plane, and the edges e_{12} , e_{13} and e_{23} are added between vertices of this embedded graph in a not-necessarily-planar way.

1.4.2 Planar plus a triangle

Another simple class of graphs that trivially has no K_6 -minors is the class of *planar-plus-triangle graphs*, constructed from a planar graph G with vertices v_1 , v_2 and v_3 by adding 3 edges incident with $\{v_1, v_2\}$, $\{v_1, v_3\}$ and $\{v_2, v_3\}$ respectively, as illustrated in Figure 1.8. This class of graphs is distinct from apex graphs because the deletion of any single vertex can remove at most two of the 3 edges in the triangle, which will leave a non-planar graph in some cases.

Proposition 1.4.2. *If G is a planar-plus-triangle graph, then G does not contain a K_6 -minor.*

Proof. Each minor of a planar-plus-triangle graph is either planar-plus-triangle or apex, but K_6 is neither planar-plus-triangle nor apex. \square




Figure 1.9: A schematic of the construction of a doublecross graph; the shaded region represents a graph embedded in the plane with vertices $u_1, u_2, v_1, v_2, u_3, u_4, v_3, v_4$ appearing in clockwise order around the outer face.

1.4.3 Doublecross

The class of *doublecross graphs* provides yet another example of a class of graphs with no K_6 -minor; a *doublecross graph* is constructed as follows: let G_0 be a graph embedded in the disk with exactly eight vertices on the boundary of the disk, labelled, in clockwise order, $u_1, u_2, v_1, v_2, u_3, u_4, v_3, v_4$; the *doublecross graph* G is constructed from G_0 by adding, for each $i \in \{1, 2, 3, 4\}$, an edge, e_i incident with vertices u_i and v_i ; see Figure 1.9.

Doublecross graphs have no K_6 -minors because it can be shown that doublecross graphs have a *linkless embedding* in 3-dimensional space, meaning that each pair of disjoint cycles are not *linked*, in the sense of knot-theory; Sachs [29] and Conway and Gordon [2] proved that K_6 has no linkless embedding and, in fact, Robertson, Seymour and Thomas [27] proved that the linklessly embeddable graphs are precisely the graphs that do not have a minor in the *Petersen family* of graphs, which includes K_6 .

Theorem 1.4.3. *If G has a linkless embedding, then G does not contain a K_6 -minor.*

Proposition 1.4.4. *If G is a doublecross graph, then G does not contain a K_6 -minor.*

1.4.4 Hose

A *hose graph* is a generalization of a doublecross graph. This class of K_6 -minor-free graphs was discovered, but not published, by Robertson, Seymour and Thomas, and appears in papers of Kawarabayashi and Mohar [15] and Kawarabayashi, Mukae and Nakamoto [16], as well as a thesis of Whalen [32], under the name *serpent*. A *hose graph* is constructed as follows:

1. let G_1, \dots, G_n be a sequence of graphs, where $n \geq 2$;
2. for $i \in \{2, \dots, n-1\}$, suppose G_i has an embedding on the disk with exactly 10 vertices on the boundary of the disk labelled, in clockwise order,

$$x_{i,1}, x'_{i,1}, x_{i,2}, x'_{i,2}, x_{i,3}, y_{i,1}, y'_{i,1}, y_{i,2}, y'_{i,2}, y_{i,3};$$

3. suppose G_1 has edges e_1 and e_2 incident with vertices $\{u_1, v_1\}$ and $\{u_2, v_2\}$, respectively, and that $G_1 - \{e_1, e_2\}$ can be embedded on the disk with exactly 9 vertices on the boundary of the disk labelled, in clockwise order,

$$u_1, u_2, v_1, v_2, y_{1,1}, y'_{1,1}, y_{1,2}, y'_{1,2}, y_{1,3};$$

4. suppose G_n has edges e_3 and e_4 incident with vertices $\{u_3, v_3\}$ and $\{u_4, v_4\}$, respectively, and that $G_n - \{e_3, e_4\}$ can be embedded on the disk with exactly 9 vertices on the boundary of the disk labelled, in clockwise order,

$$x_{n,1}, x'_{n,1}, x_{n,2}, x'_{n,2}, x_{n,3}, u_3, v_3, u_4, v_4;$$

5. for $i \in \{1, \dots, n-1\}$, let φ_i be a bijection between $\{y_{i,1}, y'_{i,1}, y_{i,2}, y'_{i,2}, y_{i,3}\}$ and $\{x_{i+1,1}, x'_{i+1,1}, x_{i+1,2}, x'_{i+1,2}, x_{i+1,3}\}$ such that, for $j \in \{1, 2, 3\}$, $\varphi_i(y_{i,j}) \in \{x_{i+1,1}, x_{i+1,2}, x_{i+1,3}\}$ and, for $j' \in \{1, 2\}$, $\varphi_i(y'_{i,j'}) \in \{x'_{i+1,1}, x'_{i+1,2}\}$;
6. the hose graph G is constructed from the disjoint union of G_1, \dots, G_n by identifying, for each $i \in \{1, \dots, n-1\}$, and each $j \in \{1, 2, 3\}$, the vertices $y_{i,j}$ and $\varphi_i(y_{i,j})$ and identifying, for each $i \in \{1, \dots, n-1\}$ and each $j' \in \{1, 2\}$, the vertices $y'_{i,j'}$ and $\varphi_i(y'_{i,j'})$;

see Figure 1.10.

Whalen [32] showed that each hose graph has a linkless embedding, which implies a hose graph cannot have a K_6 -minor.

Proposition 1.4.5. *If G is a hose graph, then G does not contain a K_6 -minor.*

1.4.5 Hamburger

One final class of K_6 -minor-free graphs is the class of *hamburger graphs*, which appear in papers of Kawarabayashi and Mohar [15] and Kawarabayashi, Mukae and Nakamoto [16],




Figure 1.10: A schematic illustration of the construction of hose graphs; the shaded regions represent the graphs G_1, \dots, G_n embedded in separate disks with the indicated vertices on the boundaries of those disks; between each of these graphs, the vertices are identified in pairs of the same colour (both white or both black); the dashed lines indicate some possible pairings that could be identified.

although neither paper includes a proof of their being K_6 -minor-free. This construction has been attributed to unpublished work of Robertson, Seymour and Thomas. A *hamburger graph* is a graph G constructed as follows: let G_1, G_2 and G_3 be three graphs; suppose that, for $i \in \{1, 2, 3\}$, G_i has an embedding in the disk with vertices $v_{i,1}, \dots, v_{i,5}$ on the boundary of the disk, in clockwise order; then the hamburger graph G is obtained from the disjoint union of G_1, G_2 and G_3 by identifying, for each $j \in \{1, \dots, 5\}$, the vertices $v_{1,j}, v_{2,j}$ and $v_{3,j}$; see Figure 1.11.

Hamburger graphs can also be shown to have linkless embeddings, so hamburger graphs do not have K_6 -minors.

Proposition 1.4.6. *If G is a hamburger graph, then G does not contain a K_6 -minor.*

Proof (Sketch). By Theorem 1.4.3, it suffices to show that G has a linkless embedding. This linkless embedding is obtained by embedding G_1, G_2 and G_3 in disjoint disks D_1, D_2 and D_3 , respectively, such that, for $i \in \{1, 2, 3\}$, G_i intersects the boundary of D_i at $\{v_{i,1}, v_{i,2}, v_{i,3}, v_{i,4}, v_{i,5}\}$, and identifying the boundaries of these three disks, so that, for $j \in \{1, 2, 3, 4, 5\}$, $v_{1,j}, v_{2,j}$ and $v_{3,j}$ are identified as a single vertex, v_j ; see Figure 1.11.

If C_1 and C_2 are two vertex-disjoint cycles in G , then it cannot be the case that both C_1 and C_2 contain more than two vertices in $\{v_1, v_2, v_3, v_4, v_5\}$, so, without loss of generality,




Figure 1.11: A schematic illustration of the construction of hamburger graphs; the shaded regions represent graphs G_1 , G_2 and G_3 embedded in separate disks, with indicated vertices on the boundary of the disk; vertices are identified along the dashed lines

$|V(C_1) \cap \{v_1, v_2, v_3, v_4, v_5\}| \leq 2$. Therefore, $E(C_1)$ can contain edges in at most two of the graphs G_1 , G_2 and G_3 . Thus, there exists a disk D such that $D \cap (D_1 \cup D_2 \cup D_3) = C_1$; hence, C_2 is disjoint from D , and, therefore, is not linked with C_1 . \square

1.4.6 A conjecture

Observe that planar-plus-triangle graphs, doublecross graphs and hamburger graphs, are not generally apex. However, deleting any two of the three vertices in the triangle turns a planar-plus triangle graph into a planar graph and deleting one vertex from each crossing pair in a doublecross graph yields a planar graph, so planar-plus-triangle graphs and doublecross graphs are 2-apex; Also, deleting any three of the five identified vertices in a hamburger graph yields a planar graph, so hamburger graphs are 3-apex.

For any natural number k , there exists a hose graph that is not k -apex, but they can be constructed using 5-sums from 4-apex graphs.

Lemma 1.4.7. *If G is a hose graph, then G can be constructed using 5-sums from 4-apex graphs.*

Proof. Let G_1, \dots, G_n be as in the definition of a hose graph. For $i \in \{1, \dots, n\}$, let G'_i be obtained from G_i by adding edges between each pair of vertices in $\{x_1, x'_1, x_2, x'_2, x_3\}$, if

$i < n$, and adding edges between each pair of vertices in $\{y_1, y'_1, y_2, y'_2, y_3\}$, if $i > 1$. Then G can be obtained from G_1, \dots, G_n by 5-sums in the obvious way.

For $i \in \{1, \dots, n\}$, deleting from G_i any two of the vertices in $\{x_1, x'_1, x_2, x'_2, x_3\}$, if $i < n$, and any two of the vertices in $\{y_1, y'_1, y_2, y'_2, y_3\}$, if $i > 1$, yields a planar graph, so G_i is 4-apex. \square

Motivated by these examples, we conjecture the following.

Conjecture 1.4.8. *There exists a natural number N such that every graph with no K_6 -minor can be obtained from graphs with at most N vertices and 4-apex graphs by clique-sums of order at most 5.*

If true, this would be exactly the sort of qualitative structure theorem for graphs with no K_6 -minor that we seek. The key improvement of Conjecture 1.4.8 over Theorem 1.2.1, is that the clique-sums used in the construction have order at most 5; as illustrated in Figure 1.4, each K_6 -minor in a graph constructed by clique-sums of order at most 5 corresponds naturally to a K_6 -minor in exactly one of the graphs used in the construction, but this is not true for graphs constructed using clique-sums of larger order. The results in this thesis are designed to help prove Conjecture 1.4.8, or something like it.

1.5 Tree-decompositions

It is convenient to describe clique-sum constructions in terms of another structure, called a *tree-decomposition*—a *tree-decomposition* of a graph G is a pair, (T, μ) , where T is a tree and μ is an injective function from the edges of G to the leaves of T . Each vertex v in $V(G)$ gives rise to a subtree, T_v , of T : the minimum tree containing the set of leaves in T that are labelled by any graph-edge incident with v . Each tree-node t in T corresponds naturally to a set of vertices in G : the vertices $v \in V(G)$ for which T_v contains t ; we call this set of vertices the *node-bag* of t . Similarly, each tree-edge f in T corresponds to the set of vertices, called the *edge-bag* of f , consisting of the vertices $v \in V(G)$ for which T_v contains f . See Figure 1.13.

Tree-decompositions have been discovered several times, by Bertelé and Brioschi [1], by Halin [13], and by Robertson and Seymour [24]. Tree-decompositions have traditionally been defined directly in terms of the node-bags, with the axioms that each pair of adjacent graph-vertices appear in a node-bag together and that, for each graph-vertex v , the set of




Figure 1.12: An example graph for which a tree-decomposition is shown in Figure 1.13




Figure 1.13: One possible tree-decomposition of the graph shown in Figure 1.12. The non-leaf tree-nodes are labelled with their node-bag; the separation presented by the bold edge is $(\{1, \dots, 16\}, \{17, \dots, 27\})$; the edge-bag of the bold edge is $\{g, i, j\}$

nodes whose node-bags containing v form a subtree. The non-standard definition here is more convenient for our purposes.

Each tree-edge f in a tree-decomposition (T, μ) of a graph G *presents* the separation (A_1, A_2) defined as follows: let T_1 and T_2 be the components of $T \setminus \{f\}$ and, for $i \in \{1, 2\}$, let A_i be the set of graph-edges $e \in E(G)$ for which $\mu(e)$ lies in T_i . Note that the edge-bag of f is precisely the set of vertices in G which are incident with an edge on each side of the separation presented by f .

Two separations (A, B) and (C, D) of a graph G *cross* if $A \cap C \neq \emptyset$, $A \cap D \neq \emptyset$, $B \cap C \neq \emptyset$ and $B \cap D \neq \emptyset$. A tree-decomposition (T, μ) of a graph G can be thought of as an encoding of a family of (pairwise) non-crossing separations as follows.

Lemma 1.5.1. *If G is a graph then*

1. *for each tree-decomposition (T, μ) of G , the separations presented by the tree-edges in T are non-crossing, and*
2. *for each family \mathcal{A} of non-crossing separations in G , there exists a unique tree-decomposition (T, μ) of G such that \mathcal{A} is precisely the family of separations presented by tree-edges in T that are not incident with a leaf.*

1.5.1 Clique-sums

Next we show how tree-decompositions are equivalent to clique-sum constructions. For each tree-decomposition (T, μ) of a graph G and each node t in T , the *part of (T, μ) at t* is the graph G_t with $V(G_t)$ equal to the node-bag of t and $E(G_t)$ defined as follows: for each tree-edge f incident with t and each pair of distinct vertices u and v in the edge-bag of f , $E(G_t)$ contains an edge $e_{f,u,v}$ incident with u and v ; if t is a leaf and $t = \mu(e)$ for some graph-edge e in G , then $E(G_t)$ additionally contains the edge e , with the same incidences as in G .

If the parts of a tree-decomposition are combined using clique-sums in the natural way, the graph G is recovered.

Theorem 1.5.2. *If G is a graph with no loops and at least 3 edges and (T, μ) is a tree-decomposition of G , then G can be constructed using clique-sums from the parts of (T, μ) .*

Proof. The proof goes by induction on the number of tree-edges f in T that are not incident with any leaf. If every tree-edge in T is incident with a leaf, then T is a star and the part of (T, μ) at the unique non-leaf node is isomorphic to G .

Suppose f is a tree-edge between non-leaf nodes t_1 and t_2 . For $i \in \{1, 2\}$, let G_i be the subgraph of G consisting of the edges, e , for which $\mu(e) \in T_i$; let G'_i be obtained from G_i by adding edges $e_{i,u,v}$ incident with vertices u and v , for each pair $\{u, v\}$ of distinct vertices in the edge-bag of f . Note that G is a clique-sum of G'_1 and G'_2 .

For $i \in \{1, 2\}$, let T_i be the component of $T \setminus \{f\}$ containing t_i and let T'_i be obtained from T_i by adding leaves $s_{i,u,v}$ adjacent to t_i for each pair $\{u, v\}$ of distinct vertices in the edge-bag of f ; for $e \in E(G_i)$, let $\mu_i(e) = \mu(e)$ and for each pair $\{u, v\}$ of distinct vertices in the edge-bag of f , let $\mu_i(e_{i,u,v}) = s_{i,u,v}$.

For $i \in \{1, 2\}$, (T'_i, μ_i) is a tree-decomposition of G'_i , and T'_i has fewer edges not incident with any leaf than T ; by the induction hypothesis, G'_i can be constructed using clique-sums from the parts of (T'_i, μ_i) , which are also parts of (T, μ) . Thus, G can be constructed using clique-sums from the parts of (T, μ) . \square

On the other hand, clique-sum constructions also give rise to tree-decompositions. To see this, note first that the set of vertices in a clique must appear together in some node-bag.

Lemma 1.5.3. *If G is a graph, (T, μ) is a tree-decomposition of G and K is a complete subgraph in G , then there exists a tree-node t_K in T such that the node-bag of t_K contains every graph-vertex in $V(K)$.*

Proof. For each vertex $v \in V(K)$, let T_v be the subtree of T consisting of the tree-nodes whose node-bag contains v . For each pair $\{u, v\}$ of vertices in $V(K)$, K contains an edge $e_{u,v}$ incident with u and v and $\mu(e_{u,v}) \in T_u \cap T_v$. Therefore, because these subtrees have pairwise non-empty intersection, there exists a node t_K in T_v for all $v \in V(K)$ (this uses the Helly property for subtrees of a tree). \square

Thus, if a graph is a clique-sum of two graphs, then tree-decompositions of the summands can be combined into a tree-decomposition of the sum. We only need to add an edge between the tree-nodes whose node-bags contain the vertices of the cliques being summed along; see Figures 1.14 and 1.15.

If G is a clique-sum of the graphs G_1 and G_2 along the cliques K_1 and K_2 in G_1 and G_2 respectively, and (T_1, μ_1) and (T_2, μ_2) are tree-decompositions of G_1 and G_2 respectively, then the *amalgamation* of (T_1, μ_1) and (T_2, μ_2) is a tree-decomposition (T, μ) , where T is constructed from the disjoint union of T_1 and T_2 by adding a single edge, f_0 , and where for each $i \in \{1, 2\}$ and each $e \in E(G_i) - E(K_i)$, $\mu(e) = \mu_i(e)$.




Figure 1.14: A clique-sum decomposition of the graph in Figure 1.12, corresponding to the bold tree-edge in Figure 1.13.

Lemma 1.5.4. *If G is a clique-sum of the graphs G_1 and G_2 along the cliques K_1 and K_2 in G_1 and G_2 respectively, and (T_1, μ_1) and (T_2, μ_2) are tree-decompositions of G_1 and G_2 respectively, and (T, μ) is the amalgamation of (T_1, μ_1) and (T_2, μ_2) , then, for $i \in \{1, 2\}$ and for each tree-edge $f \in E(T_i) \cap E(T)$, the edge-bag of f in (T, μ) is a subset of the edge-bag of f in (T_i, μ_i) , and the edge-bag of f_0 is a subset of $V(K_1) = V(K_2)$.*

This shows that a clique-sum construction of a graph gives rise to a natural tree-decomposition—each of the graphs used in the clique-sum construction has a trivial tree-decomposition which is just a star with a leaf for each edge, and, for each clique-sum operation, the tree-decompositions can be combined according to Lemma 1.5.4.

The parts of a tree-decomposition of a graph G are nearly minors of G , except that some additional edges may be needed for the clique-sums.

Lemma 1.5.5. *If G is a simple graph, (T, μ) is a tree-decomposition of G , t is a node in T , and G' is the subgraph of G induced by the node-bag of t , then G' is a subgraph of the part of (T, μ) at t , up to relabelling of edges.*

Proof. Suppose e is a graph-edge in G' incident with vertices u and v . Both u and v are in the node-bag of t , and are also in the node-bag of the leaf $\mu(e)$. If $t = \mu(e)$, then e is an edge incident with u and v in the part of (T, μ) at t . Otherwise, let f be the tree-edge

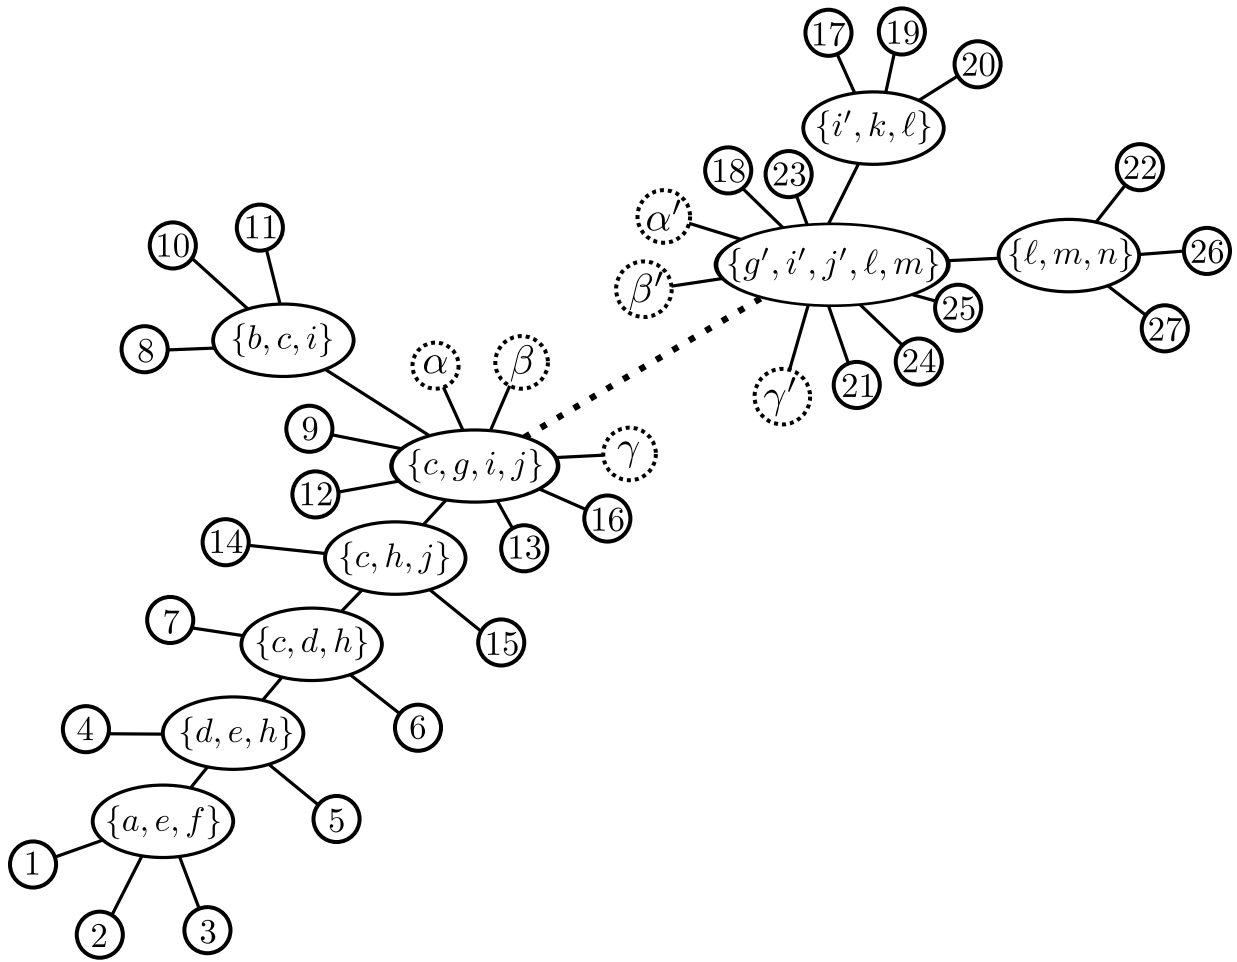


Figure 1.15: Tree-decompositions of the two clique-summands in Figure 1.14; combining these tree-decompositions by adding the dashed edge between the nodes $\{c, g, i, j\}$ and $\{g', i', j', l, m\}$, as described in Lemma 1.5.4, yields the tree-decomposition in Figure 1.13.

on the path between t and $\mu(e)$ which is incident with t . Then u and v are both in the edge-bag of f , and, hence, are adjacent in the part of (T, μ) at t . \square

1.5.2 Tree-width and branch-width

A graph is “tree-like” if it can be obtained from small graphs by clique-sums. For example, trees are highly “tree-like”, while large cliques are highly non-tree-like. Tree-decompositions give rise to two graph parameters measuring how “tree-like” a graph is.

The classic measure of “tree-like”, *tree-width*, goes back to the original work of Bertelé and Brioschi [1], Halin [13], and Robertson and Seymour [24]. For each tree-node t in a tree-decomposition, the *node-width* of t is the size of the node-bag of t . The *node-width* of a tree-decomposition (T, μ) is the maximum node-width of any tree-node in T , and the *tree-width* of a graph G , denoted $\text{tw}(G)$, is the minimum node-width of any tree-decomposition of G minus 1.

It can be shown that a simple graph G has tree-width at most 1 if and only if G is a forest, while, at the other extreme, the complete-graph K_n has tree-width $n - 1$ by Lemma 1.5.3.

For each tree-edge f in a tree-decomposition, the *edge-width* of f is the size of the edge-bag of f . The *edge-width* of a tree-decomposition (T, μ) is the maximum edge-width of any tree-edge in T .

Simply bounding the edge-width of a tree-decomposition is not enough to tell us that a graph is “tree-like” because every graph G has a tree-decomposition (T, μ) with edge-width at most 2—namely, T is the star with $|E(G)|$ leaves; the edge-bag of each tree-edge f in T consists of the two graph-vertices incident with the graph-edge $\mu^{-1}(\ell)$. The problem here, is that this tree-decomposition has an extremely high-degree node, so we need to control the degrees of the nodes in the tree-decomposition along with the edge-widths.

One natural approach, introduced by Robertson and Seymour [25], is *branch-width*—the *branch-width* of a graph G , denoted $\text{bw}(G)$, is the minimum edge-width of the tree-decompositions of G with degree at most 3.

Robertson and Seymour [25] showed that the branch-width of a graph is closely related to the tree-width, as follows. In a tree-decomposition of degree at most three, each vertex in a node-bag of a non-leaf node t is in the edge-bag of at least two of the tree-edges incident with t , so the node-width can exceed the edge-width by a factor of at most $3/2$. On the other hand, in a tree-decomposition of small node-width, high-degree nodes can be replaced by degree-3 trees, and each of the edge-bags will be a subset of some node-bag.

Theorem 1.5.6 (Robertson and Seymour). *If G is a graph with $\text{bw}(G) \geq 2$, then*

$$\text{bw}(G) \leq \text{tw}(G) + 1 \leq (3/2) \text{bw}(G).$$

If G has branch-width at most 5, then it can be constructed from small graphs (with at most 7 vertices) using clique-sums of order at most 5. This means that Conjecture 1.4.8 holds for graphs with branch-width at most 5.

1.5.3 θ -tree-width and θ -branch-degree

For Conjecture 1.4.8, we are interested only in tree-decompositions of fixed maximum edge-width, 5. For each natural number θ , a θ -tree-decomposition of G is a tree-decomposition of edge-width at most θ . For each natural number θ , the θ -tree-width of a graph G , denoted $\text{tw}_\theta(G)$, is the minimum node-width of the θ -tree-decompositions of G minus 1.

A graph with bounded 5-tree-width can be constructed from small graphs using clique-sums of order at most 5, so Conjecture 1.4.8 holds for such graphs.

Branch-width also has an analogous parameter for tree-decompositions of bounded edge-width. The θ -branch-degree of G , denoted $\text{bd}_\theta(G)$, is the minimum δ for which G has a θ -tree-decomposition of degree at most δ ; This means that graphs with bounded 5-branch-degree can be constructed from small graphs using clique-sums of order at most 5, so Conjecture 1.4.8 holds for such graphs.

The θ -tree-width and θ -branch-degree parameters are qualitatively equivalent, just as tree-width and branch-width are.

Theorem 1.5.7. *For each natural number θ , there exist positive constants c_θ and c'_θ such that if G is a graph with $\text{bd}_\theta(G) \geq 3$, then, $\text{tw}_\theta(G) \leq \text{bd}_\theta(G)\theta$ and $\text{bd}_\theta(G) \leq \binom{\text{tw}_\theta(G)}{\theta}$.*

Proof. Let (T, μ) be a θ -tree-decomposition of G such that each node in T has degree at most $\text{bd}_\theta(G)$. Then each node, t in T has at most $\text{bd}_\theta(G)\theta$ vertices in its node-bag, so $\text{tw}_\theta(G) \leq \text{bd}_\theta(G)\theta$.

Let (T', μ') be a θ -tree-decomposition of G such that the node-width of each node in T' is at most $\text{tw}_\theta(G)$ and the sum of the degrees of the nodes in T' with degree at least 4 is minimized. Let t be a tree-node in T' with degree at least 4. If f_1 and f_2 are two tree-edges incident with $\{t, s_1\}$ and $\{t, s_2\}$, respectively, and the edge-bags of f_1 and f_2 are both contained in some set $A \subseteq E(G)$ with $|A| = \theta$, then f_1 and f_2 can be replaced by a new degree-3 node, s' adjacent to t, s_1 and s_2 ; all three edges incident with s' have edge-bags

that are subsets of A , and the node-bag of s' is also a subset of A so this operation yields another θ -tree-decomposition of no larger node-width than (T', μ') , but with a smaller sum of degrees of the nodes with degree at least 4; this contradicts the choice of (T', μ') , so, in fact, the union of any two edge-bags of tree-edges incident with t must contain more than θ vertices. Thus, t has degree at most $\binom{\text{tw}_\theta(G)}{\theta}$, so $\text{bd}_\theta(G) \leq \binom{\text{tw}_\theta(G)}{\theta}$. \square

To prove Conjecture 1.4.8, it would suffice to show that a graph G has a 5-tree-decomposition (T, μ) such that the parts of (T, μ) at the nodes with large node-bags (or, equivalently, high-degree) are 4-apex.

1.5.4 Displaying a minor

Let us now examine how minors interact with separations and tree-decompositions.

First we show that separations in a graph give rise to separations in each minor of that graph.

Lemma 1.5.8. *If G is a graph, G' is a minor of G , and (A, B) is a θ -separation in G , then $(A \cap E(G'), B \cap E(G'))$ is a θ -separation in G' .*

Proof. If (A, B) is a separation in G and $e \in E(G)$, then the number of vertices incident with both an edge in A and an edge in B is at most the number of vertices incident with both an edge in $A - \{e\}$ and an edge in $B - \{e\}$, so the order of the separation $(A - \{e\}, B - \{e\})$ in $G - e$ is at most the order of the separation (A, B) in G . If z is the vertex in G/e created by contracting e , then z is incident with both an edge in $A - \{e\}$ and an edge in $B - \{e\}$, then either u or v is incident with both an edge in A and an edge in B , so the order of the separation $(A - \{e\}, B - \{e\})$ in G/e is at most the order of the separation (A, B) in G . The result follows by induction on $|E(G)| - |E(G')|$. \square

In particular, Lemma 1.5.8 shows that θ -tree-decompositions in a graph give rise to θ -tree-decompositions in each minor.

Lemma 1.5.9. *For each natural number θ , if G is a graph, (T, μ) is a θ -tree-decomposition of G and G' is a minor of G , then $(T, \mu|_{E(G)})$ is a θ -tree-decomposition of G' .*

If (T, μ) is a tree-decomposition of a graph G and t is a non-leaf node in T such that the node-bag of t is the entire set $V(G)$, then we say that t *displays* G . In general, G might be displayed by zero, one, or more nodes in T . However, if $|V(G)|$ is greater than

the edge-width of (T, μ) , then G cannot be displayed by more than one node in T and, if G is θ -connected, then G is displayed by exactly one node in any $(k - 1)$ -tree-decomposition of G .

Lemma 1.5.10. *For each natural number θ , if G is a θ -connected graph with $|V(G)| \geq \theta$ and (T, μ) is a $(\theta - 1)$ -tree-decomposition of G , then there exists a unique tree-node t in T such that t displays G .*

More generally, if G' is a minor of G , and t' is a node in T such that t' displays G' in $(T, \mu|_{E(G')})$, then we also say that t' displays G' in $(T, \mu|_{E(G)})$. By Lemma 1.5.10, each θ -connected minor of G is displayed by exactly one node in each $(\theta - 1)$ -tree-decomposition of G .

Lemma 1.5.11. *For each natural number θ , if G is a graph, G' is a θ -connected minor of G and (T, μ) is a $(\theta - 1)$ -tree-decomposition of G , then there exists a unique tree-node t in T such that t displays G' .*

In particular, K_θ is θ -connected, so each K_θ -minor of G is displayed by exactly one node in each $(\theta - 1)$ -tree-decomposition of G .

Lemma 1.5.5 implies that minors displayed at a node in a tree-decomposition are minors of the part of the tree-decomposition at that node.

Lemma 1.5.12. *If G is a graph, G' is a simple minor of G , (T, μ) is a tree-decomposition of G , and t is a tree-node in T displaying G' , then G' is isomorphic to a minor of the part of (T, μ) at t .*

With this in mind, we reformulate Wagner's Theorem in terms of tree-decompositions as follows.

Theorem 1.5.13 (Wagner's Theorem, tree-decomposition version). *If G is a graph with no K_5 -minor, then G has a 3-tree-decomposition (T, μ) such that, for each node t in T , the part of (T, μ) at t either is isomorphic to V_8 or is planar.*

This version of Wagner's Theorem is equivalent to Theorem 1.1.3 because the tree-decomposition simply encodes a cliques-sum construction. However, this formulation suggests a slightly different interpretation—the structure of the parts of the tree-decomposition certify that no K_5 -minor can be displayed at each node, by Lemma 1.5.12; Lemma 1.5.10 shows that this also certifies that the entire graph is K_5 -minor-free.

An exact structure theorem for K_6 might similarly be given in terms of a 5-tree-decomposition, together with some structure at each node t certifying that a K_6 -minor cannot be displayed at t ; for example, a 5-tree-decomposition in which each part is apex certifies that the graph is K_6 -minor-free. As we are seeking only a qualitative structure theorem, we could relax the “certificate” at each node to allow parts that are 4-apex. Of course, this would no longer certify that the graph was K_6 -minor-free, but it would still give useful structure. We can reformulate Conjecture 1.4.8 in terms of tree-decompositions as follows.

Conjecture 1.5.14. *There exists a natural number θ such that if G is a graph with no K_6 -minor then G has a 5-tree-decomposition (T, μ) such that, for each node t in T , either t has node-width at most θ or the part of (T, μ) at t is 4-apex.*

Note that Conjecture 1.5.14 implies Conjecture 1.4.8 because a graph with a 5-tree-decomposition can be constructed from the parts of that tree-decomposition by clique-sums of order at most 5, using Theorem 1.5.2.

Even if Conjecture 1.5.14 turns out to be false, it is possible that a more sophisticated qualitative structure for K_6 -minor-free graphs could be defined in terms of 5-tree-decompositions, using some alternative “relaxed certificate” at each node; such “relaxed certificates” might not even have a nice formulation in terms of the part of the 5-tree-decomposition at that node, so this formulation of the conjecture might be more robust than Conjecture 1.4.8.

1.6 Tangles

To prove something like Conjecture 1.5.14, we need to understand the conditions that can force 5-tree-decompositions to have nodes with large node-bags (or, equivalently, nodes with high degree).

To do this, we build on the notion of a *tangle*, introduced by Robertson and Seymour [25]. For each natural number θ with $\theta > 0$, a *tangle of order θ* , or simply a *θ -tangle*, in a graph G is a family, \mathcal{T} , of subsets of $E(G)$ satisfying the following axioms:

- (T1) for each $A \in \mathcal{T}$, $(A, E(G) - A)$ is a $(\theta - 1)$ -separation;
- (T2) for each $(\theta - 1)$ -separation (A, B) , either $A \in \mathcal{T}$ or $B \in \mathcal{T}$;
- (T3) for each $(\theta - 1)$ -separation (A, B) , if there exists $A' \in \mathcal{T}$ with $A \subseteq A'$, then $A \in \mathcal{T}$;

(T4) for each partition (A_1, A_2, A_3) of $E(G)$, \mathcal{T} does not contain all three of the sets A_1 , A_2 and A_3 ;

(T5) for each $e \in E(G)$, $E(G) - \{e\} \notin \mathcal{T}$.

The intuition here is that \mathcal{T} contains the “small side” of each $(\theta - 1)$ -separation, relative to a particular “ θ -connected piece” of the graph. Hence, we will refer to the sets in \mathcal{T} as \mathcal{T} -small or simply *small*. For example, if G is a large θ -connected graph, then the entire graph ought to be a “ θ -connected piece”; indeed, if $|V(G)| \geq \theta$, and (A, B) is a $(\theta - 1)$ -separation in G then, without loss of generality, each vertex in G is incident with an edge in B , but there is some vertex in G that is not incident with any edge in A ; we say that the set A is *small*, and the collection of these small sets forms a θ -tangle.

Proposition 1.6.1. *For each natural number θ with $\theta > 1$, if G is a θ -connected graph with $|V(G)| > 3(\theta - 1)$ and \mathcal{T} is the family consisting of the sets $A \subseteq E(G)$ such that $(A, E(G) - A)$ is a $(\theta - 1)$ -separation and each vertex in G is incident with at least one edge in $E(G) - A$, then \mathcal{T} is a θ -tangle in G .*

Proof. Axiom (T1) is satisfied by construction. Axiom (T2) is satisfied by the earlier observation that, by θ -connectivity of G , for each $(\theta - 1)$ -separation (A, B) , either every vertex in G is incident with an edge in A , or every vertex in G is incident with an edge in B , but not both. Axiom (T3) is satisfied because, if $A' \in \mathcal{T}$, then G contains a vertex v incident with no edge in A' , so if $A \subseteq A'$, then v contains no edge incident with A .

To prove axiom (T4), suppose (A_1, A_2, A_3) is a partition of $E(G)$ and suppose that $\{A_1, A_2, A_3\} \subseteq \mathcal{T}$. For $i \in \{1, 2, 3\}$, at most $\theta - 1$ vertices in G are incident with an edge in A_i , so at most $3(\theta - 1)$ vertices in G are incident with edges in $A_1 \cup A_2 \cup A_3$. Therefore, $A_1 \cup A_2 \cup A_3 \neq E(G)$, contradiction.

Finally, for each edge $e \in (G)$ incident with vertices u and v , G contains a vertex $x \notin \{u, v\}$, so x is not incident with any edge in $\{e\}$, so $E(G) - \{e\} \notin \mathcal{T}$ and axiom (T5) holds. \square

Tangles in minors give rise to tangles in the graph containing them; this might seem backwards, but makes sense if you think about the fact that a θ -connected minor in a graph G should be a “ θ -connected piece” of G and should give rise to a θ -tangle in G , but a minor in a θ -connected graph need not be θ -connected or give rise to any tangle.

Lemma 1.6.2. *For each natural number θ with $\theta > 0$, if G is a graph with a minor G' , \mathcal{T}' is a θ -tangle in G' , and \mathcal{T} is the family consisting of the sets $A \subseteq E(G)$ such that $(A, E(G) - A)$ is a $(\theta - 1)$ -separation and $A \cap E(G') \in \mathcal{T}'$, then \mathcal{T} is a θ -tangle in G .*

Proof. Axiom (T1) is satisfied by construction. By Lemma 1.5.8, for each $(\theta - 1)$ -separation (A, B) in G , $(A \cap E(G'), B \cap E(G'))$ is a $(\theta - 1)$ -separation in \mathcal{T} . Therefore, axioms (T2), (T3), (T4) and (T5) hold for \mathcal{T} because they hold for \mathcal{T}' . \square

We can deduce, then, one of the principal motivational examples of tangles—the tangles coming from large, highly-connected minors.

Theorem 1.6.3. *For each natural number θ with $\theta > 1$, if G is a graph with a θ -connected minor, G' , and \mathcal{T} is the family consisting of the sets $A \subseteq E(G)$ such that each vertex in G' is incident with an edge in $E(G') - A$, then \mathcal{T} is a θ -tangle in G .*

Proof. By Proposition 1.6.1 and Lemma 1.6.2. \square

1.6.1 Duality theorems

Robertson and Seymour [25] showed that tangles are precisely the structures that force a graph to have high branch-width (or, equivalently, high tree-width).

Theorem 1.6.4. *For each natural number θ with $\theta \geq 2$ and each graph G , G has branch-width at most θ if and only if G has no tangle of order greater than θ .*

For our purposes, we are interested in forcing high 5-branch-degree (or, equivalently, high 5-tree-width), rather than high branch-width. Recall that branch-width is defined in terms of tree-decompositions with maximum degree 3, while tangles cannot contain 3 sets that partition the edge set of G ; the number 3 in both of these definitions is no coincidence, and both can be generalized. The *covering-number* of a tangle \mathcal{T} in a graph G is the size of the smallest subset, \mathcal{A} , of \mathcal{T} that covers $E(G)$, meaning $\bigcup_{A \in \mathcal{A}} A = E(G)$; by definition, the covering-number of each tangle is at least 4.

The main result of Chapter 2 is the following generalization of Theorem 1.6.4:

Theorem 1.6.5. *For natural numbers θ and δ with $\theta \geq 2$ and $\delta \geq 3$, a graph G has θ -branch-degree at most δ if and only if G has no tangle of order greater than θ and covering-number greater than δ .*

1.6.2 Tree of tangles

If a graph G has small 5-branch-degree, then Conjecture 1.5.14 holds for G ; if G has high 5-branch-degree, then Theorem 1.6.5 gives a 6-tangle, \mathcal{T} , in G with high covering-number. This is some progress, but it still quite limited because the tangle is a *local* structure, representing a large “6-connected piece” of G . Other parts of G “far away from \mathcal{T} ” might have different tangles, or might have “locally small 5-branch-degree”. What we really want is a tangle for each “large 6-connected piece” of G , and a 5-tree-decomposition of small degree for the parts that are not “6-connected pieces”. Indeed, we will show that there is a 5-tree-decomposition in which each high degree-node is “justified” by a 6-tangle with high covering-number.

If

1. θ is a natural number with $\theta > 0$,
2. G is a graph,
3. \mathcal{T} is a θ -tangle in G ,
4. (T, μ) is a $(\theta - 1)$ -tree-decomposition,
5. t is a non-leaf tree-node in T ,
6. for each tree-edge f incident with t , A_f is the set of graph-edges e in G for which $\mu(e)$ and t are in distinct components of $T \setminus \{f\}$, and
7. for each tree-edge f incident with t , $A_f \in \mathcal{T}$,

then we say that the node t *displays* the tangle \mathcal{T} .

Recall from Lemma 1.5.11 that each large, θ -connected minor in G is displayed at a unique node in each $(\theta - 1)$ -tree-decomposition; θ -tangles, which generalize θ -connected minors, are also displayed at unique nodes in each $(\theta - 1)$ -tree-decomposition. This is proved in Chapter 3 as Lemma 3.2.1.

Lemma 1.6.6. *For each natural number θ with $\theta > 0$, if G is a graph, \mathcal{T} is a θ -tangle, (T, μ) is a $(\theta - 1)$ -tree-decomposition of G , then there exists a unique tree-node t in T such that t displays \mathcal{T} .*

The main result of Chapter 3 is that there always exists a θ -tree-decomposition where the degree of each node is forced by a tangle it displays.

Theorem 1.6.7. *For each natural number θ with $\theta > 1$, if G is a graph, then there exists a $(\theta - 1)$ -tree-decomposition (T, μ) of G such that, for each tree-node t in T , if $\deg_T(t) > 3$, then G contains a unique θ -tangle \mathcal{T}_t and covering-number δ_t such that t displays \mathcal{T}_t ; moreover,*

$$\delta_t \leq \deg_T(t) \leq \theta \delta_t.$$

Theorem 1.6.7 will likely be quite useful in proving a structure theorem such as Conjecture 1.5.14 because it gives a 5-tree-decomposition such that each high-degree node displays a 6-tangle with high covering-number; these tangles make the parts of this tree-decomposition at high degree nodes “weakly 6-connected”, in some sense. In this way, we are able to decompose a graph along 5-separations into “6-connected pieces”, as we alluded to earlier.

1.7 θ -connected sets

Another rich source of tangles comes from the θ -connected set of vertices in a graph—for $\theta \in \mathbb{N}$, a set of vertices, X in a graph G is said to be a θ -connected set if, for each pair of subsets $Y, Z \subseteq X$ with $|Y| = |Z| \leq \theta$, G contains a collection of $|Y|$ vertex-disjoint paths between Y and Z .

A natural example a θ -connected set arises in the $\theta \times \theta$ -grid—the graph with vertex set $\{v_{i,j} : i, j \in \{1, \dots, \theta\}\}$ with $v_{i,j}$ adjacent to $v_{i',j'}$ if and only if $|i - i'| + |j - j'| = 1$. For each $i \in \{1, \dots, \theta\}$, the i^{th} -row of G , $\{v_{i,j} : j \in \{1, \dots, \theta\}\}$, is a θ -connected set; see Figure 1.16.

A large, θ -connected set in a graph gives rise to a tangle in the following natural way.

Theorem 1.7.1. *For natural numbers θ and δ with $\theta \geq 1$ and $\delta \geq 3$, if X is a θ -connected set in a graph G with $|X| > (\theta - 1)(\delta - 1)$, and \mathcal{T} is the family consisting of the sets $A \subseteq E(G)$ such that $(A, E(G) - A)$ is a $(\theta - 1)$ -separation in G and fewer than θ of vertices in X are incident with an edge in A , then \mathcal{T} is a θ -tangle with covering number at least δ .*

Conversely, each θ -tangle with high covering-number also give rise to a large θ -connected set, although in a more complicated way that we defer to Chapter 2.

Theorem 1.7.2. *For positive integers θ and n there exists a positive integer δ such that, if G is a graph and \mathcal{T} is a θ -tangle in G with covering number at least δ , then G contains a θ -connected set of size at least n .*




Figure 1.16: An example of a collection of vertex-disjoint paths between the vertex sets Y and Z in a single row of the grid; in general, similar collections paths exist of any pair of sets of vertices in a single row of a $\theta \times \theta$ -grid, so the set of vertices in each row of a $\theta \times \theta$ -grid form an θ -connected set.

This correspondence between large, θ -connected sets and θ -tangles with high covering number, together with the correspondence between tree-width and branch-width in Theorem 1.5.6, shows that the existence of large, highly-connected sets in a graph G is qualitatively equivalent to G having high tree-width. There is also more direct connection between large, highly connected sets and high tree-width, shown by Diestel, Jensen, Gorbonov and Thomassen [4].

Theorem 1.7.3. *For each natural number ω , and for each graph G ,*

- (i) *if G contains an $(\omega + 1)$ -connected set of size at least 3ω then G has tree-width at least ω , and*
- (ii) *conversely, if G has no $(\omega + 1)$ -connected set of size at least 3ω , then G has tree-width less than 4ω .*

1.8 Minors arising from θ -tangles

Robertson and Seymour's classic Grid Theorem [23] (see also [4]) shows that high-order tangles give rise to large grid-minors.

Theorem 1.8.1 (Grid Theorem). *For each $n \in \mathbb{N}$, there exists some $N \in \mathbb{N}$ such that if G has a tangle of order N , then G contains an $n \times n$ -grid-minor.*

The Grid Theorem provides an important dual for branch-width (or, equivalently, tree-width)—by Theorem 1.6.4 and Theorem 1.8.1, a graph either has large branch-width or it has a large grid-minor. A grid-minor in a graph is generally easier to make use of in applications than a tangle.

1.8.1 The main theorem

For studying graphs with no K_6 -minor, the Grid Theorem cannot be applied because we are not interested in high-order tangles, but instead in 6-tangles, with high covering-number. We generalize the Grid Theorem by showing that θ -tangles with high covering-number give rise to a graph constructed as follows. For each $r, \ell, n \in \mathbb{N}$ with $r \geq 1$ and $n \geq 3$, an (r, ℓ, n) -wheel is a graph G constructed by taking the union of the following pieces, illustrated in Figure 1.17:

1. for $i \in \{1, \dots, n\}$, let T_i be an r -vertex tree such that, for distinct $i, j \in \{1, \dots, n\}$, T_i and T_j are vertex-disjoint;
2. for $i \in \{1, \dots, n\}$, let M_i be a perfect matching between $V(T_i)$ and $V(T_{i+1})$ (indices taken modulo n);
3. let Z be a set of vertices with $|Z| = \ell$ such that Z is disjoint from $V(T_i)$ for each $i \in \{1, \dots, n\}$;
4. for $i \in \{1, \dots, n\}$ and $z \in Z$, there is a unique edge $e_{i,z}$ incident with z and a vertex in $V(T_i)$.

We show in Chapter 4 that a θ -tangle with high covering-number gives rise to a minor isomorphic to either the complete bipartite graph, $K_{\theta,n}$ or an (r, ℓ, n) -wheel, where $2r + \ell = \theta$ and n is large. This is the main result of this thesis.

Theorem 1.8.2. *For natural numbers θ, n with $\theta \geq 2$, there exists a natural number m such that if G is a graph containing a θ -tangle with covering-number at least m , then G contains either a $K_{\theta,n}$ -minor, or an (r, ℓ, n) -wheel-minor such that $r, \ell \in \mathbb{N}$ and $2r + \ell = \theta$.*

Moreover, we show that, if $2r + \ell$ and m are both sufficiently large, then each (r, ℓ, m) -wheel contains an $n \times n$ -grid-minor, implying the Grid Theorem.




Figure 1.17: A (5, 3, 10)-wheel.

1.8.2 Specialization for $\theta = 6$

For $\theta = 6$, Theorem 1.8.2 can be refined to show that each graph with a 6-tangle of sufficiently high covering-number contains a minor from one of the following families, each depicted in Figure 1.18.

- (a) $K_{6,n}$ —the complete bipartite graph with 6 vertices on one side and n vertices on the other side
- (b) a cycle of length n plus 4 vertices each adjacent to every vertex in the cycle
- (c),(d) a cyclic ladder of length n , plus two vertices each adjacent to all vertices on one side of the cyclic ladder
- (e),(f) a Möbius (twisted) ladder of length n , plus two vertices each adjacent to all vertices on one side of the Möbius ladder
- (g),(h),(i),(j) a cyclic triple-ladder of length n , with each possible twist.

We hope that a proof of Conjecture 1.5.14 can be found by understanding the ways in which graphs in Figure 1.18 can be extended to non-4-apex graphs. Norine and Thomas [20] give an exact description of minimal non-planar extensions of planar graphs, and also give a description of non- θ -apex extensions of θ -apex graphs in some limited settings, so the prospects of such research seem promising.

1.9 Summary of results

This thesis presents new tools for developing qualitative structure theorems for classes of graphs that exclude small minors, such as K_6 . Our main results are summarized below.

1. We prove Theorem 2.6.1, showing an exact duality between $(\theta - 1)$ -branch-degree and the covering number of a θ -tangle. This result generalizes the duality between tangles and branch-width proved by Robertson and Seymour [25].
2. We prove Theorem 3.1.1, showing that each graph admits a $(\theta - 1)$ -tree-decomposition where the degree of each node is either 3 or is qualitatively related to the covering number of the unique θ -tangle displayed at the node. This result can be interpreted as a global decomposition of a graph into its “ θ -connected components”.




Figure 1.18: The unavoidable-minors in tangles of order 6, each representing an infinite family.

3. Finally, we prove Theorem 4.1.7, showing that each graph that contains a sufficiently large θ -connected set of vertices contains one of a finite set of explicitly described minors, each of which has a large θ -connected set. This is the main result of the thesis; it implies the Grid Theorem (Theorem 4.1.6) of Robertson and Seymour [23].

Chapter 2

Tree-decompositions and tangles

2.1 Introduction

Robertson and Seymour [25] introduced tangles (see also [10]) and proved that they are dual to branch-width:

Theorem 2.1.1. *For $\theta \in \mathbb{N}$ with $\theta \geq 2$, if G is a graph, then G has branch-width at most θ if and only if G has no tangle of order greater than θ .*

The main goal of this chapter is to prove Corollary 2.6.2 which shows that θ -branch-degree and $(\theta + 1)$ -tangles are dual in a similar same way.

There is also a natural *qualitative* duality between tree-width and large, highly-connected sets, due to Diestel, Jensen, Gorbonov and Thomassen [4].

Theorem 2.1.2. *For each natural number ω , and for each graph G ,*

- (i) if G contains an $(\omega + 1)$ -connected set of size at least 3ω then G has tree-width at least ω , and*
- (ii) conversely, if G has no $(\omega + 1)$ -connected set of size at least 3ω , then G has tree-width less than 4ω .*

Together with the qualitative equivalence of branch-width and tree-width, these results show that high-order tangles and large, highly-connected sets of vertices are qualitatively

equivalent. We generalize this and give a direct proof that θ -tangles with high covering-number are qualitatively equivalent to large θ -connected sets in Theorem 2.9.2.

In summary, this chapter proves that the following are qualitatively equivalent for each natural number $\theta > 2$ and each graph G :

- (1) G has high $(\theta - 1)$ -tree-width;
- (2) G has high $(\theta - 1)$ -branch-degree;
- (3) G has a θ -tangle with high covering-number;
- (4) G has a large θ -connected set.

The qualitative equivalence between $(\theta - 1)$ -tree-width (1) and $(\theta - 1)$ -branch-degree (2) is established in Theorem 1.5.7. Corollary 2.6.2 shows that the $(\theta - 1)$ -branch-degree of G (2) equals the maximum covering-number of a θ -tangle of G (3). Theorem 2.9.2 shows that the presence of a θ -tangle with high covering-number (3) implies the presence of a large, θ -connected set (4); conversely, Theorem 2.9.1 shows that the presence of a large, θ -connected set (4) implies the presence of a θ -tangle with high covering-number (3).

Like Robertson and Seymour [25], we obtain our results in the more general setting of “connectivity systems”.

2.2 Connectivity systems

A *connectivity system* is a pair (S, λ) consisting of a *ground set* S and a *connectivity function* $\lambda : 2^S \rightarrow \mathbb{N}$ such that

- λ is *symmetric*; that is $\lambda(X) = \lambda(S - X)$ for each $X \subseteq S$; and
- λ is *submodular*; that is $\lambda(X \cup Y) + \lambda(X \cap Y) \leq \lambda(X) + \lambda(Y)$ for each $X, Y \subseteq S$.

A *separation* in a connectivity system is a bipartition of the ground set. The *order* of a separation (X, Y) is $\lambda(X)$; for each natural number θ , a separation of order at most θ is called a θ -separation. Connectivity systems were introduced by Fujishige [9], where he showed that each connectivity system has a canonical decomposition along 2-separations.

2.2.1 Examples of connectivity systems

We are principally interested in connectivity systems that arise from graphs in the following way. For each graph G , there is a connectivity system whose ground set is $E(G)$ and whose connectivity function, λ_G , is defined for $F \subseteq E(G)$ as the number of vertices incident with both an edge in F and an edge in $E(G) - F$. Fujishige [9] prove that this is a connectivity system.

Lemma 2.2.1. *If G is a graph then $(E(G), \lambda_G)$ is a connectivity system.*

Connectivity systems arising from graphs in this way are called a *graphic connectivity system*.

Connectivity systems also arise naturally from matroids. Given a matroid M with ground set E and rank function ρ , there is a connectivity system with ground set E and connectivity function λ_M defined for $X \subseteq E$ as $\lambda_M(X) = \rho(X) + \rho(E - X) - \rho(E)$; λ_M is symmetric by definition and submodular because ρ is submodular.

2.2.2 Domination

In this section we explore a natural substructure relation for connectivity systems, which is related to the concept of a minor in a graph or a matroid.

If (S, λ) is a connectivity system, then, for any two disjoint sets A and B in S , the *connectivity between A and B* is

$$\kappa_\lambda(A, B) = \min\{\lambda(A') : A \subseteq A' \subseteq S - B\}.$$

In the case of a graphic connectivity system of a graph G , $\kappa_{\lambda_G}(A, B)$ measures the maximum number of vertex-disjoint paths containing both a vertex incident with an edge in A and a vertex incident with an edge in B .

If (S, λ) and (S', λ') are connectivity systems such that $S' \subseteq S$ and, for each $X \subseteq S'$, $\lambda'(X) \leq \kappa_\lambda(X, S' - X)$, then we say that (S, λ) *dominates* (S', λ') .

Lemma 2.2.2. *If (S, λ) and (S', λ') are connectivity systems and (S, λ) dominates (S', λ') , then, for each $A \subseteq S$, $\lambda(A) \geq \lambda'(A \cap S')$.*

Proof.

$$\lambda(A \cap S') \leq \kappa_\lambda(A \cap S', S' - A) \leq \lambda(A). \quad \square$$

The principal example of dominated connectivity systems arise from minors in a graph.

Lemma 2.2.3. *If G is a graph and G' is a minor of G , then $(E(G), \lambda_G)$ dominates $(E(G'), \lambda_{G'})$.*

Proof. Suppose C and D are disjoint subsets of edges of G and $G' = G/C \setminus D$; that is, G' is obtained from G by contracting the edges in C and deleting the edges in D .

Let $F' \subseteq E(G')$ and choose $F \subseteq E(G)$ such that $F' \subseteq F \subseteq E(G') - F'$ and $\lambda_G(F) = \kappa_{\lambda_G}(F', E(G') - F')$. Let U' be the set of vertices, u , in G' incident with both an edge, $e_{1,u}$ in F' and an edge, $e_{2,u}$ in $E(G') - F'$. For each $u \in U'$, there exists a path P_u in G such that $E(P_u) \subseteq C$ and P_u contains a vertex incident with $e_{1,u}$ and a vertex incident with $e_{2,u}$; moreover, for distinct $u, v \in U'$, P_u and P_v are vertex-disjoint. Thus, if $F \subseteq E(G)$ such that $F' \subseteq F \subseteq E(G') - F'$, and $\lambda_G(F) = \kappa_{\lambda_G}(F', E(G') - F')$ then

$$\kappa_{\lambda_G}(F', E(G') - F') = \lambda_G(F) \geq |U'| = \lambda_{G'}(F'),$$

so (G, λ_G) dominates $(G', \lambda_{G'})$. \square

2.3 Tree-decompositions

We describe here how the concept of tree-decompositions for graphs, as defined in Chapter 1, generalizes to connectivity systems.

A *leaf* in a tree T is a node of degree exactly 1. An *incidence* in a tree T is a pair (t, e) , where t is a node and e is a tree-edge incident with t in T . A *leaf-incidence* in a tree T is an incidence (t, e) such that t is a leaf.

Tree-decompositions, discovered independently by Bertelé and Brioschi [1], by Halin [13], and by Robertson and Seymour [24] in the context of graphs, generalize in a trivial way to connectivity systems—a *tree-decomposition* of a connectivity system (S, λ) is a pair (T, μ) where T is a tree and μ is an injective function from the ground set, S , to the set of leaves of T .

If (T, μ) is a tree-decomposition of the connectivity system (S, λ) , then each tree-edge e in T gives rise to a separation in (S, λ) as follows. Note that $T \setminus \{e\}$ has two components,

T_1 and T_2 . For $i \in \{1, 2\}$, let $S_i \subseteq S$ be the elements, s , of the ground set for which $\mu(s)$ is in T_i . We say that (S_1, S_2) is the *separation presented by e* in (T, μ) . If t_1 is the node incident with e in T with $t_1 \in V(T_1)$, then we also say that S_2 is the *set presented by the incidence (t_1, e)* ; we emphasize that the set presented by (t_1, e) is the set of elements labelling leaves across e from t_1 , *not* the set of elements labelling leaves on the same side of e as t_1 .

If (S, λ) is a connectivity system with a tree-decomposition (T, μ) and e is a tree-edge in T presenting the separation (S_1, S_2) , then the *edge-width* of e is $\lambda(S_1)$ (which equals $\lambda(S_2)$ by symmetry of λ); The *edge-width* of (T, μ) is the maximum edge-width of the tree-edges in T . For $\theta \in \mathbb{N}$, the θ -*branch-degree* of a connectivity system (S, λ) , denoted $\text{bd}_k(S, \lambda)$, is the minimum $\delta \in \mathbb{N}$ for which (S, λ) has a θ -tree-decomposition with maximum degree at most δ .

2.4 Tangles

Tangles—structures which capture the general notion of a “ θ -connected piece” of a graph—were discovered by Robertson and Seymour [25]. For each natural number θ with $\theta > 0$, a *tangle of order θ* , or simply a θ -*tangle*, in a connectivity system (S, λ) is a family, \mathcal{T} , of subsets of S satisfying the following axioms:

- (T1) if $A \in \mathcal{T}$ then $\lambda(A) < \theta$;
- (T2) if $A \subseteq S$ with $\lambda(A) < \theta$ then either $A \in \mathcal{T}$ or $S - A \in \mathcal{T}$;
- (T3) if $A \subseteq A' \subseteq S$, $A' \in \mathcal{T}$ and $\lambda(A) < \theta$, then $A \in \mathcal{T}$;
- (T4) for each partition (A_1, A_2, A_3) of S , \mathcal{T} does not contain all three of the sets A_1 , A_2 and A_3 ;
- (T5) for each $x \in S$, $S - \{x\} \notin \mathcal{T}$.

Traditionally, tangles are defined as families satisfying (T1), (T2), (T5) and the following axiom:

- (T4') if $A_1, A_2, A_3 \in \mathcal{T}$, then $A_1 \cup A_2 \cup A_3 \neq S$.

These two definitions are equivalent by the following lemma, which follows from (2.9) in [25].

Lemma 2.4.1. *For $\theta \in \mathbb{N}$, if (S, λ) is a connectivity system, and $\mathcal{T} \subseteq 2^S$ satisfying axioms (T1), (T2) and (T5) for being a θ -tangle, then \mathcal{T} is a θ -tangle if and only if \mathcal{T} satisfies (T4').*

As we saw in Proposition 1.6.1, if G is a θ -connected graph with more than $3(\theta - 1)$ vertices, then the graphic connectivity system $(E(G), \lambda_G)$ has a θ -tangle, \mathcal{T} , such that, for each $A \in \mathcal{T}$, there exists a vertex in G not incident with any edge in A .

Grids also possess natural tangles; recall that, for each natural number n , the $n \times n$ -grid is the graph with vertices $\{v_{i,j} : i, j \in \{1, \dots, n\}\}$ where $v_{i,j}$ is adjacent to $v_{i',j'}$ if and only if $|i - i'| + |j - j'| = 1$; for $i \in \{1, \dots, n\}$, the i^{th} -row of the $n \times n$ -grid is the subgraph induced by the vertices $\{v_{i,j} : j \in \{1, \dots, n\}\}$; The following is attributed to Kleitman and Saks [25].

Proposition 2.4.2. *For each natural number θ with $\theta \geq 3$, if G is the $\theta \times \theta$ -grid, and $\mathcal{T} \subseteq 2^{E(G)}$ is the family consisting of the sets $A \subseteq E(G)$ for which there exists $i_A \in \{1, \dots, \theta\}$ such that each vertex in the i_A^{th} -row of G is not incident with any edge in A , then \mathcal{T} is a θ -tangle in G .*

2.4.1 Induced tangles

As we saw in Section 1.6, a tangle in a graph G' gives rise to a tangle in each graph G containing G' as a minor, a tangle in a connectivity system (S', λ') gives rise to a tangle in each connectivity system (S, λ) dominating (S', λ') .

Lemma 2.4.3. *For each natural number θ , if (S, λ) and (S', λ') are connectivity systems such that (S, λ) dominates (S', λ') , \mathcal{T}' is a θ -tangle in (S', λ') , and $\mathcal{T} \subseteq 2^S$ is the family consisting of the sets $A \subseteq S$ for which $A \cap S' \in \mathcal{T}'$, then \mathcal{T} is a θ -tangle in (S, λ) .*

Proof. Let \mathcal{T} be the family consisting of the sets $A \subseteq S$ such that $\lambda(A) < \theta$ and $A \cap S' \in \mathcal{T}'$. Then \mathcal{T} satisfies (T1) by construction and satisfies (T2), ..., (T5) because \mathcal{T}' satisfies these axioms, together with the fact that, for each $(\theta - 1)$ -separation (A, B) in (S, λ) , $(A \cap S', B \cap S')$ is a $(\theta - 1)$ -separation in (S', λ') by Lemma 2.2.2. \square

We say that the tangle \mathcal{T} in Lemma 2.4.3 is *induced* by the tangle \mathcal{T}' .

If a graph G has a large, θ -connected minor, then G has a θ -tangle induced by the tangle in the θ -connected minor; similarly, if G has a $\theta \times \theta$ -grid-minor, then G has a θ -tangle induced by the tangle in the $\theta \times \theta$ -grid.

Induced tangles can also be used to relate tangles with sets of vertices, using the following connectivity system: if G is graph and $(E(G), \lambda_G)$ is the graphic connectivity system for G , then $(E(G), \lambda_G)$ is dominated by the connectivity system $(E(G) \cup V(G), \tilde{\lambda}_G)$, where, for each $F \subseteq E(G)$ and each $U \subseteq V(G)$, $\tilde{\lambda}_G(F \cup U)$ is the number of vertices, v in G such that both

1. $v \in U$ or v is incident with an edge in F , and
2. $v \in V(G) - U$ or v is incident with an edge in $E(G) - F$.

Note that $(E(G) \cup V(G), \tilde{\lambda}_G)$ is a connectivity system and dominates $(E(G), \lambda_G)$ because $(E(G) \cup V(G), \tilde{\lambda}_G)$ is isomorphic to the graphic connectivity system on the graph \tilde{G} obtained from G by adding a loop edge at each vertex, and G is a minor of \tilde{G} . We call $(E(G) \cup V(G), \tilde{\lambda}_G)$ the *extended graphic connectivity system* of G . By Lemma 2.4.3, each tangle in $(E(G), \lambda_G)$ induces a tangle in $(E(G) \cup V(G), \tilde{\lambda}_G)$ of the same order.

2.5 Covering-number

The *covering-number* of a tangle \mathcal{T} in a connectivity system (S, λ) is the size of the smallest subset, \mathcal{A} , of \mathcal{T} such that $\bigcup_{A \in \mathcal{A}} A = S$.

One natural class of graphs with tangles of high covering-number are the *cylindrical grids*—the $\theta \times n$ -*cylindrical-grid* is the graph with vertex set $\{v_{i,j} : i \in \{1, \dots, \theta\}, j \in \mathbb{Z}_n\}$ (where \mathbb{Z}_n denotes the cyclic group of order n) such that vertices $v_{i,j}$ and $v_{i',j'}$ are adjacent if and only if

$$(i - i', j - j') \in \{(0, 1), (0, -1), (1, 0), (-1, 0)\};$$

see Figure 2.1. For $i \in \{1, \dots, \theta\}$, the i^{th} -*circuit* of the $\theta \times n$ -cylindrical-grid is the subgraph induced by the vertex set $\{v_{i,j} : j \in \mathbb{Z}_n\}$; for $j \in \mathbb{Z}_n$, the j^{th} -*path* of the $\theta \times n$ -cylindrical-grid is the subgraph induced by the vertex set $\{v_{i,j} : i \in \{1, \dots, \theta\}\}$.

Proposition 2.5.1. *For $\theta, \delta, n \in \mathbb{N}$ with $\theta \geq 2$ and $n > (\delta - 1)(2\theta - 1)$, the graphic connectivity system of the $\theta \times n$ -cylindrical-grid contains a 2θ -tangle \mathcal{T} with covering number at least δ .*

Proof. Let G be the $\theta \times n$ -cylindrical-grid. Let \mathcal{T} be the family consisting of the sets $A \subseteq E(G)$ such that $\lambda_G(A) < 2\theta$ and there exists $i_A \in \{1, \dots, \theta\}$ for which A does not contain any edge in the i_A^{th} -circuit of G . Then \mathcal{T} satisfies axioms (T1), (T3) and (T5) by construction.




Figure 2.1: A 5×20 -cylindrical grid.

To prove axiom (T2), let (A, B) be a $(2\theta - 1)$ -separation in (S, λ) . For, $i \in \{1, \dots, \theta\}$, if the i^{th} -circuit of G contains both an edge in A and an edge in B , then the i^{th} -circuit contains two vertices that are each incident with both an edge in A and an edge in B ; hence, this occurs for at most $\theta - 1$ circuits so, without loss of generality, there exists $i_A \in \{1, \dots, \theta\}$ such that the edges in the i_A^{th} -circuit are each in B . Thus, $A \in \mathcal{T}$ and \mathcal{T} satisfies axiom (T2).

To prove axiom (T4), note that, for each $A \in \mathcal{T}$ and each $j \in \mathbb{Z}_n$, if the j^{th} -path of G contains an edge in A , then, because $v_{i_A, j}$ is incident with an edge in $E(G) - A$, the j^{th} -path contains a vertex incident in G with both an edge in A and an edge in $E(G) - A$; therefore, there are at most $(2\theta - 1)$ such paths. Thus, if $A_1, \dots, A_{\delta-1} \in \mathcal{T}$, then at most $(\delta - 1)(2\theta - 1)$ paths of G contain an edge in $\bigcup_{i=1}^{\delta-1} A_i$, so $\bigcup_{i=1}^{\delta-1} A_i \neq E(G)$. Hence, axiom (T4) holds and \mathcal{T} has covering-number at least δ . \square

Cylindrical grids are a special case of the class of graphs that we describe in Chapter 4, which also have tangles with high covering number.

We observe that the covering-number of a θ -tangle gives a lower bound on the maximum degree of a $(\theta - 1)$ -tree-decompositions. This gives the first (easy) direction of the duality between θ -tangles with high covering-number and $(\theta - 1)$ -tree-decompositions of low degree.

Lemma 2.5.2. *For each natural number θ , if (S, λ) is a connectivity system, \mathcal{T} is a θ -tangle in (S, λ) and (T, μ) is a $(\theta - 1)$ -tree-decomposition of (S, λ) , then T has a node with degree at least the covering-number of \mathcal{T} .*

Proof. Let \vec{T} be an orientation of T obtained as follows. If t is a node in T incident with a tree-edge e and the set presented by the incidence (t, e) is in \mathcal{T} , then orient e away from t . Note that this orientation is well defined because the separation presented by each tree-edge is a $(\theta - 1)$ -separation, and so exactly one side of the separation is in \mathcal{T} . Because \vec{T} is acyclic, there exists a node, t_0 such that each edge incident with t_0 is oriented toward t_0 . Note that t_0 is not a leaf by axiom (T5). Therefore, if \mathcal{A} is the family consisting of, for each tree-edge e incident with t_0 , the set presented by the incidence (t_0, e) , then $\mathcal{A} \subseteq \mathcal{T}$ and $\bigcup_{A \in \mathcal{A}} A = S$. Therefore, t_0 has degree at least the covering-number of \mathcal{T} . \square

2.6 The duality theorem

The main theorem of this chapter is that the converse of Lemma 2.5.2 also holds—that if (S, λ) has no θ -tangle with covering number at least δ , then (S, λ) has a $(\theta - 1)$ -tree-decomposition with maximum degree less than δ . This generalizes Robertson and Seymour’s [25] duality theorem between high order tangles and branch-width (Theorem 2.1.1).

Theorem 2.6.1. *For natural numbers θ and δ with $\delta \geq 4$, if (S, λ) is a connectivity-system with $\lambda(\{x\}) < \theta$ for each $x \in S$, and (S, λ) has no θ -tangle with covering at least δ , then (S, λ) has a $(\theta - 1)$ -tree-decomposition with maximum degree less than δ .*

Together, Lemma 2.5.2 and Theorem 2.6.1 imply the following.

Corollary 2.6.2. *For each natural number θ and each connectivity system (S, λ) , if (T, μ) is a minimum degree $(\theta - 1)$ -tree-decomposition of (S, λ) and \mathcal{T} is a θ -tangle in (S, λ) with maximum covering-number, then the degree of T equals the covering-number of \mathcal{T} .*

For each tree T , let $L(T)$ denote the set of leaves of T .

A *partial tree-decomposition* of a connectivity system (S, λ) is a pair (T, μ) where T is a tree and $\mu : S \rightarrow L(T)$. If $\mathcal{A} \subseteq 2^S$ and, for each leaf $t \in L(T)$, there exists a set $A \in \mathcal{A}$ such that $\{x \in S : \mu(x) = t\} \subseteq A$, then we say that (T, μ) is a *partial tree-decomposition over \mathcal{A}* . Note that a partial tree-decomposition over $\{\{x\} : x \in S\}$ is a tree-decomposition.

The *separation presented by a tree-edge*, the *set presented by a tree-incidence*, the *width of a tree-edge*, and the *width of a partial tree-decomposition* are all defined for partial tree-decompositions in exactly the same way that they are defined for tree-decompositions. As for tree-decompositions, for $\theta \in \mathbb{N}$, we define a *partial θ -tree-decomposition* to be a partial tree-decomposition of width at most θ .

Lemma 2.6.3. For natural numbers θ and δ with $\delta \geq 4$, if (S, λ) is a connectivity system and $\mathcal{A} \subseteq 2^S$ such that

- (1) $\bigcup_{A \in \mathcal{A}} A = S$,
- (2) for $A \in \mathcal{A}$, $\lambda(A) < \theta$,
- (3) (S, λ) has no partial $(\theta - 1)$ -tree-decomposition over \mathcal{A} of degree less than δ , and
- (4) for each $X \subseteq S$ with $\lambda(X) < \theta$, there exists $A \in \mathcal{A}$ such that either $X \subseteq A$ or $S - X \subseteq A$,

then \mathcal{A} extends to a θ -tangle in (S, λ) with covering number at least δ .

Proof. Let \mathcal{T} be the family consisting of the sets $X \subseteq S$ for which $\lambda(X) < \theta$ and $X \subseteq A$ for some $A \in \mathcal{A}$.

The family \mathcal{T} satisfies tangle property (T1) by construction. Tangle property (T2) is satisfied by assumption (4).

Claim 2.6.3.1. If $\{X_1, \dots, X_{\delta-1}\} \subseteq \mathcal{T}$, then $\bigcup_{i=1}^{\delta-1} X_i \neq S$.

Proof of Claim. Suppose for contradiction that $\{X_1, \dots, X_{\delta-1}\} \subseteq \mathcal{T}$, and $\bigcup_{i=1}^{\delta-1} X_i = S$. Choose such $\{X_1, \dots, X_{\delta-1}\}$ minimizing $\sum_{i=1}^{\delta-1} |X_i|$.

Note that $\{X_1, \dots, X_{\delta-1}\}$ is a partition of S —if $X_i \cap X_j \neq \emptyset$ for distinct i and j , then, by submodularity, either $\lambda(X_i - X_j) \leq \lambda(X_i)$ or $\lambda(X_j - X_i) \leq \lambda(X_j)$; without loss of generality, the former occurs, so $X_i - X_j \in \mathcal{T}$, contradicting minimality of $\sum_{i=1}^{\delta-1} |X_i|$.

Let T be the degree- $(\delta - 1)$ star with leaves $\{t_1, \dots, t_{\delta-1}\}$ and let $\mu : S \rightarrow L(T)$ be defined as $\mu(x) = t_i$ for each $x \in X_i$. Then (T, μ) is a partial $(\theta - 1)$ -tree-decomposition over \mathcal{A} , contradicting (4). □ (Claim)

Therefore, \mathcal{T} satisfies property (T4') and has covering-number at least δ . If $S - \{x\} \in \mathcal{T}$ for some $x \in S$, then there exists $A \in \mathcal{A} \subseteq \mathcal{T}$ such that $x \in A$, and $(S - \{x\}) \cup A = S$, which contradicts Claim 2.6.3.1; therefore, \mathcal{T} satisfies property (T5), and is a θ -tangle. □

Lemma 2.6.4. For natural numbers θ and δ with $\delta \geq 4$, if (S, λ) is a connectivity system and $\mathcal{A} \subseteq 2^S$ such that

- (1) $\bigcup_{A \in \mathcal{A}} A = S$,

(2) for $A \in \mathcal{A}$, $\lambda(A) < \theta$, and

(3) (S, λ) has no partial $(\theta - 1)$ -tree-decomposition over \mathcal{A} of degree less than δ ,

then \mathcal{A} extends to a θ -tangle in (S, λ) with covering number at least δ .

Proof. Suppose the contrary and let $\mathcal{A} \subseteq 2^S$ be maximal such that \mathcal{A} satisfies (1), (2) and (3) but \mathcal{A} does not extend to a θ -tangle with covering number at least δ .

If, for each $X \subseteq S$ with $\lambda(X) < \theta$, there exists $A \in \mathcal{A}$ such that either $X \subseteq A$ or $S - X \subseteq A$, then \mathcal{A} extends to a θ -tangle with covering-number at least δ by Lemma 2.6.3.

Suppose, then, that there exists a $(\theta - 1)$ -separation (X_1, X_2) in (S, λ) such that, for each $A \in \mathcal{A}$, $X_1 \not\subseteq A$ and $X_2 \not\subseteq A$; choose such a separation minimizing $\lambda(X_1)$.

For $i \in \{1, 2\}$, $\mathcal{A} \cup \{X_i\}$ satisfies (1) and (2), and $\mathcal{A} \cup \{X_i\}$ does not extend to a θ -tangle, so (S, λ) has a partial $(\theta - 1)$ -tree-decomposition (T_i, μ'_i) over $\mathcal{A} \cup \{X_i\}$ of degree less than δ . By assumption (T_i, μ'_i) is not a partial tree-decomposition over \mathcal{A} , so there exists some $t_i \in V(T_i)$ such that, for each $A \in \mathcal{A}$,

$$\{x \in S : \mu'_i(x) = t_i\} \not\subseteq A.$$

Because (T_i, μ'_i) is a partial tree-decomposition over $\mathcal{A} \cup \{X_i\}$,

$$\{x \in S : \mu'_i(x) = t_i\} \subseteq X_i.$$

Let $\mu_i : S \rightarrow L(T)$ be defined as

$$\mu_i(x) = \begin{cases} t_i & x \in X_i \\ \mu'_i & x \notin X_i. \end{cases}$$

Claim 2.6.4.1. For each $i \in \{1, 2\}$ and each tree-edge $e \in E(T_i)$, the width of e in (T_i, μ_i) is at most the width of e in (T_i, μ'_i)

Proof of Claim. Let t be the node incident with e such that t and t_i are in distinct components of $T_i \setminus \{e\}$. Let Y and Y' be the sets presented by the incidence (t, e) in (T_i, μ_i) and (T_i, μ'_i) respectively. Note that $Y = Y' \cup X_i$. Suppose for contradiction that $\lambda(Y) > \lambda(Y')$.

By submodularity, $\lambda(Y' \cap X_i) < \lambda(X_i)$. Therefore, by minimality of $\lambda(X_1)$, there is some $A \in \mathcal{A}$ such that either

(i) $Y' \cap X_i \subseteq A$,

(ii) or $S - (Y' \cap X_i) \subseteq A$.

In case (i),

$$\{x \in S : \mu'_i(x) = t_i\} \subseteq Y' \cap X_i \subseteq A$$

contradicting our choice of t_i . In case (ii),

$$S - X_i \subseteq S - (Y' \cap X_i) \subseteq A,$$

contradicting our choice of (X_1, X_2) . □ (Claim)

Therefore, (T_i, μ_i) is a partial $(\theta - 1)$ -tree-decomposition.

Note that, if $S = \emptyset$, then the empty tree gives a trivial partial $(\theta - 1)$ -tree-decomposition over \mathcal{A} , so we may assume $S \neq \emptyset$.

Claim 2.6.4.2. *For each $i \in \{1, 2\}$ and each leaf $t \in L(T_i) - \{t_i\}$, there exists $A \in \mathcal{A}$ such that*

$$\{x \in S : \mu_i(x) = t\} \subseteq A$$

Proof of Claim. If

$$\{x \in S : \mu_i(x) = t\} = \emptyset,$$

then, because $S \neq \emptyset$, we know that $\mathcal{A} \neq \emptyset$, so $\emptyset \subseteq A$ for an arbitrary $A \in \mathcal{A}$.

Otherwise, there is some $x \in S$ such that $\mu_i(x) = t \neq t_i$, so

$$\{x \in S : \mu_i(x) = t\} \not\subseteq X_i$$

but, because (T_i, μ'_i) is a partial tree-decomposition over $\mathcal{A} \cup \{X_i\}$, there exists $A \in \mathcal{A}$ such that

$$\{x \in S : \mu_i(x) = t\} \subseteq \{x \in S : \mu'_i(x) = t\} \subseteq A.$$

□ (Claim)

Let T be the tree constructed the disjoint union of T_1 and T_2 by identifying nodes t_1 and t_2 . Let $\mu : S \rightarrow L(T)$ be defined for $x \in S$ as

$$\mu(x) = \begin{cases} \mu_1(x) & x \in X_2 \\ \mu_2(x) & x \in X_1. \end{cases}$$

Note that $L(T) = (L(T_1) \cup L(T_2)) - \{t_1, t_2\}$, so, by Claim 2.6.4.2, μ is well-defined and (T, μ) is a partial tree-decomposition over \mathcal{A} . For $i \in \{1, 2\}$, each tree-edge in T_i presents the same separation in (T_i, μ_i) and (T, μ) , so, by Claim 2.6.4.1, (T, μ) is a partial $(\theta - 1)$ -tree-decomposition over \mathcal{A} . This contradicts our choice of \mathcal{A} . \square

We now prove Theorem 2.6.1, restated here.

Theorem 2.6.1. *For natural numbers θ and δ with $\delta \geq 4$, if (S, λ) is a connectivity-system with $\lambda(\{x\}) < \theta$ for each $x \in S$, and (S, λ) has no θ -tangle with covering at least δ , then (S, λ) has a $(\theta - 1)$ -tree-decomposition with maximum degree less than δ .*

Proof. Let $\mathcal{A} = \{\{x\} : x \in S\}$. If (S, λ) has no $(\theta - 1)$ -tree-decomposition, then it has no partial $(\theta - 1)$ -tree-decomposition over \mathcal{A} . Therefore, by Lemma 2.6.4, \mathcal{A} extends to a θ -tangle in (S, λ) . \square

2.7 Tangle matroids

Each tangle in a connectivity system gives rise to a *matroid*; here we use standard definition of a matroid in terms of the *rank function*; that is, a *matroid* is a pair (S, ρ) where S is a set, called the *ground set* of the matroid, and $\rho : 2^S \rightarrow \mathbb{N}$, called the *rank function*, such that

(M1) for $A \subseteq S$, $\rho(A) \leq |A|$,

(M2) ρ is *monotonic*—for $A, B \subseteq S$, if $A \subseteq B$ then $\rho(A) \leq \rho(B)$, and

(M3) ρ is *submodular*—for $A, B \subseteq S$, $\rho(A \cup B) + \rho(A \cap B) \leq \rho(A) + \rho(B)$.

Given a connectivity system (S, λ) and a θ -tangle \mathcal{T} , the *rank-function of \mathcal{T}* , $\rho_{\mathcal{T}} : 2^S \rightarrow \mathbb{N}$ is defined for $X \subseteq S$ as

$$\rho_{\mathcal{T}}(X) = \begin{cases} \min\{\lambda(A) : X \subseteq A \in \mathcal{T}\} & \text{if } X \subseteq A \text{ for some } A \in \mathcal{T} \\ \theta & \text{otherwise.} \end{cases}$$

Intuitively, $\rho_{\mathcal{T}}(X)$ gives the “connectivity of X into the tangle \mathcal{T} ”. We have called $\rho_{\mathcal{T}}$ the “rank-function” of \mathcal{T} because we would like it to be the rank-function of a matroid. Indeed, $(S, \rho_{\mathcal{T}})$ satisfies matroid axioms (M2) and (M3), as shown in Lemma 2.7.2 and

Lemma 2.7.3, respectively, but matroid axiom (M1) might not hold in general. For example, there might be some $x \in S$ such that, for each $A \in \mathcal{T}$ with $x \in A$, $\lambda(A) > 1$, so $\rho_{\mathcal{T}}(\{x\}) > 1 = |\{x\}|$. Robertson, Seymour and Thomas [28] (see also [10]) proved that restricting the connectivity of the empty-set and the singletons yields a matroid.

Lemma 2.7.1. *For $\theta \in \mathbb{N}$ with $\theta \geq 2$, if (S, λ) is a connectivity system with $\lambda(\emptyset) = 0$ and, for each $x \in S$, $\lambda(\{x\}) \leq 1$, \mathcal{T} is a θ -tangle in (S, λ) and $\rho_{\mathcal{T}}$ is the rank-function of \mathcal{T} , then $(S, \rho_{\mathcal{T}})$ is a rank- θ matroid.*

Here we want something slightly more general, although the proof is not substantially different. For $V \subseteq S$ with $\lambda(\{v\}) \leq 1$ for each $v \in V$, we define the *tangle-matroid of \mathcal{T} on V* , $M(V, \mathcal{T}) = (V, \rho')$, where ρ' is the restriction of $\rho_{\mathcal{T}}$ to V . We show that that $M(V, \mathcal{T})$ is a matroid in Lemma 2.7.5, using the following lemmas.

Lemma 2.7.2 and Lemma 2.7.3 are proved by Robertson, Seymour and Thomas [28].

Lemma 2.7.2. *For $\theta \in \mathbb{N}$ with $\theta \geq 2$, if (S, λ) is a connectivity system with $\lambda(\emptyset) = 0$, \mathcal{T} is a θ -tangle in (S, λ) , then the rank-function of \mathcal{T} is monotonic.*

Proof. Suppose $X \subseteq Y \subseteq S$. If $\rho_{\mathcal{T}}(Y) = \theta$, then $\rho_{\mathcal{T}}(X) \leq \theta = \rho_{\mathcal{T}}(Y)$. We may assume, then, that $\rho_{\mathcal{T}}(Y) < \theta$, so there exists $A \in \mathcal{T}$ such that $Y \subseteq A$ and $\rho_{\mathcal{T}}(Y) = \lambda(A)$. Therefore, $X \subseteq A$, so $\rho_{\mathcal{T}}(X) \leq \lambda(A) = \rho_{\mathcal{T}}(Y)$. \square

Lemma 2.7.3. *For $\theta \in \mathbb{N}$ with $\theta \geq 2$, if (S, λ) is a connectivity system with $\lambda(\emptyset) = 0$, \mathcal{T} is a θ -tangle in (S, λ) , then the rank-function of \mathcal{T} is submodular.*

Proof. Suppose first that, for each $A \in \mathcal{T}$, $X \not\subseteq A$. Then for each $A \in \mathcal{T}$, $X \cup Y \not\subseteq A$, so $\rho_{\mathcal{T}}(X) = \rho_{\mathcal{T}}(X \cup Y) = \theta$. Thus,

$$\rho_{\mathcal{T}}(X \cup Y) + \rho_{\mathcal{T}}(X \cap Y) = \rho_{\mathcal{T}}(X) + \rho_{\mathcal{T}}(X \cap Y) \leq \rho_{\mathcal{T}}(X) + \rho_{\mathcal{T}}(Y),$$

where the last inequality holds by Lemma 2.7.2, because $X \cap Y \subseteq Y$.

If, for each $A \in \mathcal{T}$, $Y \not\subseteq A$, then $\rho_{\mathcal{T}}(X \cup Y) + \rho_{\mathcal{T}}(X \cap Y) \leq \rho_{\mathcal{T}}(X) + \rho_{\mathcal{T}}(Y)$ by the same argument.

We may assume, then, that there exists $A, B \in \mathcal{T}$ such that $X \subseteq A$, $Y \subseteq B$, $\rho_{\mathcal{T}}(X) = \lambda(A)$ and $\rho_{\mathcal{T}}(Y) = \lambda(B)$.

Note that $\rho_{\mathcal{T}}(X \cup Y) \leq \lambda(A \cup B)$; indeed, if $A \cup B \in \mathcal{T}$, then, because $X \cup Y \subseteq A \cup B$, $\rho_{\mathcal{T}}(X \cup Y) \leq \lambda(A \cup B)$; on the other hand, if $A \cup B \notin \mathcal{T}$, then, by tangle axiom (T4), $S - (A \cup B) \notin \mathcal{T}$, so, by tangle axiom (T2), $\lambda(A \cup B) \geq \theta$, and $\rho_{\mathcal{T}}(X \cup Y) \leq \theta \leq \lambda(A \cup B)$.

Similarly, $\rho_{\mathcal{T}}(X \cap Y) \leq \lambda(A \cap B)$. Thus,

$$\begin{aligned} \rho_{\mathcal{T}}(X \cup Y) + \rho_{\mathcal{T}}(X \cap Y) &\leq \lambda(A \cup B) + \lambda(A \cap B) \\ &\leq \lambda(A) + \lambda(B) \\ &= \rho_{\mathcal{T}}(X) + \rho_{\mathcal{T}}(Y). \quad \square \end{aligned}$$

Lemma 2.7.4. *For $\theta \in \mathbb{N}$ with $\theta \geq 2$, if (S, λ) is a connectivity system with $\lambda(\emptyset) = 0$, \mathcal{T} is a θ -tangle in (S, λ) , and $X \subseteq S$ with $|X| < \theta$ and $\lambda(\{x\}) \leq 1$ for each $x \in X$, then $X \in \mathcal{T}$ and $\lambda(X) \leq |X|$.*

Proof. By induction on $|X|$.

Note that $\lambda(\emptyset) = 0 \leq \theta$, so $\emptyset \in \mathcal{T}$ and $\lambda(\emptyset) = 0$.

Note also that, for each $v \in V$, $\lambda(\{v\}) \leq 1 \leq \theta$, so $\{v\} \in \mathcal{T}$ and $\lambda(\{v\}) = 1$.

We may assume, then, that $1 < |X| < \theta$ and, for each $Y \subseteq V$, if $|Y| < |X|$ then $Y \in \mathcal{T}$ and $\lambda(Y) \leq |Y|$. Let $x \in X$.

$$\begin{aligned} \lambda(X) &= \lambda(X) + \lambda(\emptyset) \\ &\leq \lambda(X - \{x\}) + \lambda(\{x\}) \\ &\leq (|X| - 1) + 1 = |X|. \end{aligned}$$

Therefore, $\lambda(X) < \theta$. By the induction hypothesis, $X - \{x\}$ and $\{x\}$ are both in \mathcal{T} , so, by tangle axiom (T4), $S - X \notin \mathcal{T}$, so by tangle axiom (T2), $X \in \mathcal{T}$. \square

Now we can prove that $M(V, \mathcal{T})$ is a matroid.

Lemma 2.7.5. *For $\theta \in \mathbb{N}$ with $\theta \geq 2$, if (S, λ) is a connectivity system with $\lambda(\emptyset) = 0$, \mathcal{T} is a θ -tangle in (S, λ) , and $V \subseteq S$ with $\lambda(\{v\}) \leq 1$ for each $v \in V$, then $M(V, \mathcal{T})$ is a matroid.*

Proof. Let ρ' be the restriction of $\rho_{\mathcal{T}}$ to V .

To prove that $M(V, \mathcal{T})$ satisfies matroid axiom (M1), note that, for $X \subseteq V$ with $|X| \geq \theta$, $\rho_{\mathcal{T}}(X) \leq \theta \leq |X|$; therefore, it suffices to show that, for $X \subseteq V$, with $|X| < \theta$, $\rho_{\mathcal{T}}(X) \leq |X|$; by Lemma 2.7.4, if $|X| < \theta$, then $X \in \mathcal{T}$, so $\rho_{\mathcal{T}}(X) \leq \lambda(X) \leq |X|$.

$M(V, \mathcal{T})$ satisfies matroid axiom (M2) by Lemma 2.7.2, and satisfies matroid axiom (M3) by Lemma 2.7.3. \square

2.8 Graphic tangles and their matroids

For graphic connectivity systems, tangle matroids allow us to relate tangles with the vertices of the graph. Suppose G is a graph and \mathcal{T} is a θ -tangle in $(E(G), \lambda_G)$. As we saw in Section 2.4, the graphic connectivity system $(E(G), \lambda_G)$ is dominated by the extended graphic connectivity system $(E(G) \cup V(G), \tilde{\lambda}_G)$, and the tangle \mathcal{T} induces a tangle, $\tilde{\mathcal{T}}$ in $(E(G) \cup V(G), \tilde{\lambda}_G)$. The *graphic tangle-matroid*, $M(G, \mathcal{T})$ is the tangle-matroid $M(V(G), \tilde{\mathcal{T}})$ in $(E(G) \cup V(G), \tilde{\lambda}_G)$. We let $\rho_{\mathcal{T}}(X)$ denote the rank-function of $M(G, \mathcal{T})$. Then, for each $X \subseteq V(G)$, $\rho_{\tilde{\mathcal{T}}}(X)$ essentially gives “connectivity of X into \mathcal{T} ”.

Lemma 2.8.1. *For $\theta, k \in \mathbb{N}$ with $\theta > 2$, if G is a graph, \mathcal{T} is a θ -tangle in G , and $X_1, X_2 \subseteq V(G)$, then G contains $\min(\rho_{\mathcal{T}}(X_1), \rho_{\mathcal{T}}(X_2))$ vertex-disjoint paths between X_1 and X_2 .*

Proof. Suppose for contradiction that G does not contain $\min\{\rho_{\mathcal{T}}(X_1), \rho_{\mathcal{T}}(X_2)\}$ vertex-disjoint paths between X_1 and X_2 . By Menger’s Theorem (see [3, Theorem 3.3.1]), there exists graphs G_1 and G_2 such that

- (a) $G = G_1 \cup G_2$,
- (b) $X_1 \subseteq V(G_1)$,
- (c) $X_2 \subseteq V(G_2)$, and
- (d) $|V(G_1) \cap V(G_2)| < \min\{\rho_{\mathcal{T}}(X_1), \rho_{\mathcal{T}}(X_2)\}$.

Then, for $i \in \{1, 2\}$,

$$\lambda_G(E(G_i)) \leq \tilde{\lambda}_G(E(G_i) \cup V(G_i)) \leq |V(G_1) \cap V(G_2)| < \min\{\rho_{\mathcal{T}}(X_1), \rho_{\mathcal{T}}(X_2)\} \leq \theta.$$

We may assume that $E(G_1) \cap E(G_2) = \emptyset$. By tangle axiom (T2), without loss of generality, $E(G_1) \in \mathcal{T}$; therefore, $E(G_1) \cup V(G_1) \in \tilde{\mathcal{T}}$, where $\tilde{\mathcal{T}}$ is the tangle in $(E(G) \cup V(G), \tilde{\lambda}_G)$ induced by \mathcal{T} . Thus, $\rho_{\mathcal{T}}(X_1) \leq \tilde{\lambda}_G(E(G_1) \cup V(G_1))$, contradicting the earlier observation that $\tilde{\lambda}_G(E(G_1) \cup V(G_1)) < \rho_{\mathcal{T}}(X_1)$. \square

2.9 θ -connected sets

Recall from Chapter 1 that, for $\theta \in \mathbb{N}$, a θ -connected set in a graph G is a set $X \subseteq V(G)$ such that, for $Y, Z \subseteq X$ with $|Y| = |Z| \leq \theta$, G contains θ vertex-disjoint paths between Y and Z .

Each large, θ -connected set in a graph gives rise to a tangle.

Theorem 2.9.1. *For $\theta, \delta \in \mathbb{N}$, with $\theta > 2$ and $\delta > 3$, if G is a graph and X is a θ -connected set in G with $|X| > (\theta - 1)(\delta - 1)$, and $\mathcal{T} \subseteq 2^{E(G)}$ is the family consisting of the sets $A \subseteq E(G)$ for which both $\lambda_G(A) < \theta$ and X contains at most $\lambda_G(A)$ vertices incident with edges in A , then \mathcal{T} is a θ -tangle in G with covering-number at least δ .*

Proof. Note that \mathcal{T} satisfies (T1) by construction. If (A, B) is a θ -separation in G then G contains at most $\lambda_G(A)$ vertex-disjoint paths between the set of vertices incident with any edge in A and the set of vertices incident with any edge in B ; because X is θ -connected, one of these sets of vertices must have size at most $\lambda_G(A)$, so either $A \in \mathcal{T}$ or $B \in \mathcal{T}$, and hence (T2) holds. Axiom (T5) holds because, for each $e \in E(G)$, the number of the vertices in X incident with an edge in $E(G) - \{e\}$, is at least

$$|X| - 2 \geq (\theta - 1)(\delta - 1) - 1 \geq 5 > \lambda(\{e\}).$$

If $A_1, \dots, A_{\delta-1} \in \mathcal{T}$ then the number of vertices in X incident with an edge in $\bigcup_{i=1}^{\delta-1} A_i$ is at most $(\delta - 1)(\theta - 1) < |X|$. But $|X| > 2$, so each vertex in X is incident with at least one edge, so $\bigcup_{i=1}^{\delta-1} A_i \neq E(G)$. Thus, (T4') holds and \mathcal{T} is a θ -tangle with covering number at least δ . \square

We will prove in this section that, qualitatively, the converse of Theorem 2.9.1 also holds.

Theorem 2.9.2. *For $\theta, n \in \mathbb{N}$, with $\theta > 2$ there exists $N \in \mathbb{N}$ such that, if G is a graph and \mathcal{T} is a θ -tangle in $(V(G), \lambda_G)$ with covering-number at least N , then G contains a θ -connected set X of size n such that, for each $A \in \mathcal{T}$, at most $\lambda_G(A)$ vertices in X are incident with any edge in A .*

Combining Theorem 2.9.1 and Theorem 2.9.2, we obtain the following qualitative equivalence between θ -connected sets and θ -tangles.

Corollary 2.9.3. *For $\theta, \delta \in \mathbb{N}$, with $\theta > 2$ and $\delta > 3$, there exist $N_{\theta, \delta}, M_{\theta, \delta} \in \mathbb{N}$ such that, for each graph G :*

1. *if G contains a θ -connected set of size at least $N_{\theta, \delta}$, then G contains a θ -tangle with covering-number at least δ , and*
2. *conversely, if G contains a θ -tangle with covering-number at least $M_{\theta, \delta}$, then G contains a θ -connected set of size at least δ .*

2.9.1 Uniform matroids

If (S, ρ) is a matroid and $X \subseteq S$, then the *restriction* of (S, ρ) to X is the matroid $(X, \rho|_X)$.

For $\theta, n \in \mathbb{N}$, the *rank- θ uniform matroid with n elements*, denoted $U_{\rho, n}$ is the matroid with ground set $\{1, \dots, n\}$ and rank function ρ_θ defined for $X \subseteq \{1, \dots, n\}$ as

$$\rho_\theta(X) = \begin{cases} |X| & \text{if } |X| \leq \theta \\ \theta & \text{if } |X| > \theta. \end{cases}$$

Robertson, Seymour and Thomas [28] showed that the independent sets in a tangle matroid are fully connected to each other. This implies that uniform restrictions of tangle matroids give rise to θ -connected sets, as follows.

Lemma 2.9.4. *For $\theta \in \mathbb{N}$ with $\theta > 2$, if*

1. G is a graph,
2. \mathcal{T} is a θ -tangle in $(E(G), \lambda_G)$,
3. $\tilde{\mathcal{T}}$ is the tangle induced by \mathcal{T} in $(E(G) \cup V(G), \tilde{\lambda}_G)$,
4. $X \subseteq V(G)$, and
5. $(X, \rho_{\tilde{\mathcal{T}}, V(G)}|_X)$, the tangle matroid of $\tilde{\mathcal{T}}$ on $V(G)$, restricted to X , is isomorphic to a rank- θ uniform matroid,

then X is a θ -connected set and, for each $A \in \mathcal{T}$, at most $\lambda(A)$ vertices in X are incident with any edge in A .

Proof. For $Y, Z \subseteq X$ with $|Y| = |Z| \leq \theta$, $\rho_{\tilde{\mathcal{T}}, V(G)}(Y) = \rho_{\tilde{\mathcal{T}}, V(G)}(Z) = |Y|$, so, by Lemma 2.8.1, G contains $|Y|$ vertex-disjoint paths between Y and Z . \square

2.9.2 Hyperplanes in matroids

In this section we prove two elementary results from matroid theory that we require.

A *hyperplane* in a matroid (S, ρ) is a maximal set $H \subseteq S$ with $\rho(H) = \rho(S) - 1$. Recall that each $X \subseteq S$ extends to a unique maximal set of rank $\rho(X)$, called the *closure* of X .

If many hyperplanes are required to cover the ground set of a matroid, then Geelen and Kabel [12] that the matroid has a large uniform restriction of rank $\rho(S)$.

Lemma 2.9.5. *For $\theta, n \in \mathbb{N}$, if (S, ρ) is a matroid with $\rho(S) = \theta$ and, S is not the union of fewer than $\binom{n-1}{\theta-1} + 1$ hyperplanes in (S, ρ) , then (S, ρ) has a restriction $(X, \rho|_X)$ isomorphic to $U_{\theta, n}$.*

2.9.3 Constructing a θ -connected set

The covering-number of a tangle is precisely the number of hyperplanes required to cover the ground set of the tangle matroid, so a tangle with high covering-number has a large uniform restriction.

Lemma 2.9.6. *For $\theta, \delta, n \in \mathbb{N}$ with $\delta > \binom{n-1}{\theta-1}$, if (S, λ) is a connectivity system, \mathcal{T} is a θ -tangle in (S, λ) and $V \subseteq S$ such that V is not covered by fewer than δ sets in \mathcal{T} , then the tangle matroid $(V, \rho_{\mathcal{T}, V})$ has a restriction isomorphic to $U_{\theta, n}$.*

Proof. Note that $\rho_{\mathcal{T}, V}(V) = \theta$. Therefore, if H is a hyperplane in $(V, \rho_{\mathcal{T}, V})$, then $\rho_{\mathcal{T}, V}(H) < \theta$, so there exists $A_H \in \mathcal{T}$ such that $H \subseteq A_H$. Hence, the number of hyperplanes required to cover V in the matroid $(V, \rho_{\mathcal{T}, V})$ is at least δ , the number of sets in \mathcal{T} required to cover V . By Lemma 2.9.5, $(V, \rho_{\mathcal{T}, V})$ has a $U_{\theta, n}$ -restriction, as desired. \square

Now we are able to construct the θ -connected set arising from a θ -tangle.

Theorem 2.9.2. *For $\theta, n \in \mathbb{N}$, with $\theta > 2$ there exists $N \in \mathbb{N}$ such that, if G is a graph and \mathcal{T} is a θ -tangle in $(V(G), \lambda_G)$ with covering-number at least N , then G contains a θ -connected set X of size n such that, for each $A \in \mathcal{T}$, at most $\lambda_G(A)$ vertices in X are incident with any edge in A .*

Proof. Let

$$N = \binom{n-1}{\theta-1} + \binom{(\theta-1)\binom{n-1}{\theta-1}}{2} + 1.$$

Let $\tilde{\mathcal{T}}$ be the tangle induced by \mathcal{T} in the extended graphic connectivity system $(E(G) \cup V(G), \tilde{\lambda}_G)$. We emphasize that the sets in $\tilde{\mathcal{T}}$ are subsets of $E(G) \cup V(G)$.

Claim 2.9.6.1. *$V(G)$ is not covered by any collection of $\binom{n-1}{\theta-1}$ sets in $\tilde{\mathcal{T}}$.*

Proof of Claim. Suppose $\mathcal{A} \subseteq \tilde{\mathcal{T}}$ such that $V(G) \subseteq \bigcup_{A \in \mathcal{A}} A$ and $|\mathcal{A}| \leq \binom{n-1}{\theta-1}$. Let $V' \subseteq V(G)$ be the set of vertices, v , in G which are incident with an edge, e_v , in $E(G) - \bigcup_{A \in \mathcal{A}} A$.

Note that, for each $v \in V'$, there is some $A_v \in \mathcal{A}$ such that $v \in A_v$ and $e_v \notin A_v$. For each $A \in \mathcal{A}$,

$$|\{v \in V' : A = A_v\}| \leq \tilde{\lambda}(A_v) < \theta.$$

Thus,

$$|V'| \leq (\theta - 1)|\tilde{\mathcal{A}}| \leq (\theta - 1) \binom{n-1}{\theta-1}.$$

For each pair $u, v \in V'$, let $E_{u,v}$ be the set of edges in G incident only with the vertices u and v ; note that $\tilde{\lambda}(E_{u,v}) \leq 2 < \theta$, so $E_{u,v} \in \mathcal{T}$. Thus,

$$\mathcal{A}' = \{A \cap E(G) : A \in \mathcal{A}\} \cup \{E_{u,v} : u, v \in V'\}$$

is a cover of $E(G)$ by sets in \mathcal{T} and

$$|\mathcal{A}'| \leq \binom{n-1}{\theta-1} + \binom{(\theta-1)\binom{n-1}{\theta-1}}{2} < N.$$

This contradicts the fact that \mathcal{T} has covering number at least N . □ (Claim)

Let $(V(G), \rho_{\tilde{\mathcal{T}}, V(G)})$ be the tangle matroid of $\tilde{\mathcal{T}}$ over $V(G)$. By Lemma 2.9.6, there exists $X \subseteq V(G)$ such that $|X| = n$ and the restriction $(X, \rho_{\tilde{\mathcal{T}}, V(G)}|_X)$ is isomorphic to $U_{\theta, n}$. By Lemma 2.9.4, X is a θ -connected set.

For $A \in \mathcal{T}$, if $X_A \subseteq X$ is the set of vertices in X incident with at least one edge in A , then $X_A \cup A \in \tilde{\mathcal{T}}$. Therefore,

$$\rho_{\tilde{\mathcal{T}}, V(G)}(X_A) \leq \tilde{\lambda}(X_A \cup A) \leq \lambda(A) < \theta$$

so, because the restriction of $(S, \rho_{\tilde{\mathcal{T}}, V(G)})$ to X is rank- θ -uniform,

$$|X_A| = \rho_{\tilde{\mathcal{T}}, V(G)}(X_A) \leq \lambda(A).$$

Thus, at most $\lambda(A)$ vertices in X are incident with any edge in A , as desired. □

Chapter 3

Trees of tangles

3.1 Introduction

The goal of Chapter 2 was to find a $(\theta - 1)$ -tree-decomposition with node-degrees as small as possible. We saw that θ -tangles with high covering number are obstacles to finding such decompositions. We will show that more is true—by Lemma 3.2.1, if (T, μ) is a $(\theta - 1)$ -tree-decomposition and \mathcal{T} is a θ -tangle, then \mathcal{T} is displayed by a unique node of T ; moreover, the degree of that node is at least the covering number of \mathcal{T} .

The goal of this chapter is to find a width $(\theta - 1)$ -tree-decomposition with node-degrees as small as possible, subject to the obstacles posed by the θ -tangles. In particular, we prove the following result.

Theorem 3.1.1. *For $\theta \in \mathbb{N}$ with $\theta > 2$, if G is a graph, then there exists a $(\theta - 1)$ -tree-decomposition (T, μ) of G such that for each node t in T with $\deg_T(t) > 3$, G contains a unique θ -tangle \mathcal{T}_t such that \mathcal{T}_t is displayed by t ; moreover, if δ_t is the covering number of \mathcal{T}_t , then*

$$\deg_T(t) \leq \theta \delta_t.$$

This result can be viewed as a global structure theorem, decomposing a graph into its “ θ -connected pieces”.

We prove the results in the same general setting of *connectivity systems* used in Chapter 2.

3.2 Displaying tangles

For $\theta \in \mathbb{N}$, if (S, λ) is a connectivity system with a θ -tangle \mathcal{T} and $(\theta - 1)$ -tree-decomposition (T, μ) , and t is a node in T such that, for each tree-edge e incident with t , the set presented by the incidence (t, e) is in \mathcal{T} , then we say that t *displays* \mathcal{T} .

Let $I(T)$ denote the set of *incidences* in a tree T ; that is, the set of pairs (t, e) where t is a node in T and e is a tree-edge incident with t .

Each θ -tangle is displayed by exactly one node of each $(\theta - 1)$ -tree-decomposition:

Lemma 3.2.1. *If (S, λ) is a connectivity system with a θ -tangle \mathcal{T} and a $(\theta - 1)$ -tree-decomposition (T, μ) , then there exists a unique tree-node $t_{\mathcal{T}}$ in T such that $t_{\mathcal{T}}$ displays \mathcal{T} ; moreover, for each incidence $(t, e) \in I(T)$, (t, e) presents a set in \mathcal{T} if and only if t and $t_{\mathcal{T}}$ are in the same component of $T \setminus \{e\}$.*

Proof. Let \vec{T} be the orientation of T such that, for each incidence (t, e) presenting a set in \mathcal{T} , e is oriented toward t ; this is a well-defined orientation because exactly one side of the $(\theta - 1)$ -separation presented by each tree-edge e is in \mathcal{T} . Note that a node, t , in T displays \mathcal{T} if and only if each tree-edge incident with t is oriented toward t in \vec{T} ; such a node, $t_{\mathcal{T}}$, exists because \vec{T} is acyclic.

Claim 3.2.1.1. *If t is a node incident with tree-edges e_1 and e_2 in T , then at least one of e_1 or e_2 is oriented toward t .*

Proof of Claim. Suppose e_1 and e_2 are both oriented away from t in \vec{T} and X_1 and X_2 are the sets presented by the incidences (t, e_1) and (t, e_2) respectively. Then $S - X_1$ and $S - X_2$ are both in \mathcal{T} . However, X_1 and X_2 are disjoint so

$$(S - X_1) \cup (S - X_2) = S,$$

which contradicts tangle property (T3). □ (Claim)

If t is a node in T and P is the (unique) path in T from t to $t_{\mathcal{T}}$ then P contains an edge oriented toward $t_{\mathcal{T}}$ (the edge in P incident with $t_{\mathcal{T}}$), so, by Claim 3.2.1.1, each edge of P is oriented toward $t_{\mathcal{T}}$. Therefore, the only tree-edge incident with t that is oriented away from t is the edge in P . Thus, for each tree-edge e incident with t , (t, e) displays a separation in \mathcal{T} if and only if t and $t_{\mathcal{T}}$ are in the same component of $T \setminus \{e\}$. This also shows that $t_{\mathcal{T}}$ is the unique node displaying \mathcal{T} . □

We say that the node $t_{\mathcal{T}}$ in Lemma 3.2.1 is *the* node in (T, μ) displaying the tangle \mathcal{T} .

The covering-number of a tangle gives a trivial lower bound on the degree of the node displaying that tangle:

Lemma 3.2.2. *If (S, λ) is a connectivity system with a θ -tangle \mathcal{T} with covering-number δ and a $(\theta - 1)$ -tree-decomposition (T, μ) , and $t_{\mathcal{T}}$ is the node displaying \mathcal{T} in (T, μ) , then $\deg_T(t_{\mathcal{T}}) \geq \delta$.*

Proof. If (t, e) is a leaf-incident in T , then (t, e) presents either the set S (if t is not in the image of μ), or the set $S - \{x\}$, for some $x \in S$ (and $\mu(x) = t$). By tangle property (T5), $S - \{x\} \notin \mathcal{T}$ and, by tangle property (T4), $S \notin \mathcal{T}$. Hence, $t_{\mathcal{T}}$ is not a leaf.

For each tree-edge e incident with $t_{\mathcal{T}}$, let X_e be the set presented by the incidence $(t_{\mathcal{T}}, e)$. Because $t_{\mathcal{T}}$ displays \mathcal{T} , $X_e \in \mathcal{T}$ for each $e \in \mathcal{T}$. Also, for each $x \in S$, if e is the tree-edge separating $t_{\mathcal{T}}$ from $\mu(x)$, then $x \in X_e$, so

$$\bigcup_{e \text{ incident with } t} X_e = S.$$

Therefore, t must be incident with at least δ tree-edges. □

3.3 Distinguishing separations

Suppose \mathcal{T}_1 and \mathcal{T}_2 are two distinct θ -tangles in a connectivity system (S, λ) . A $(\mathcal{T}_1, \mathcal{T}_2)$ -*distinguishing separation* is a $(\theta - 1)$ -separation (A_1, A_2) such that $A_1 \in \mathcal{T}_1$ and $A_2 \in \mathcal{T}_2$. Robertson and Seymour [25] proved that any two distinct tangles of the same order have a distinguishing separation.

Lemma 3.3.1. *For $\theta \in \mathbb{N}$, if \mathcal{T}_1 and \mathcal{T}_2 are distinct θ -tangles in a connectivity system (S, λ) , then there exists a $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation.*

Proof. Because \mathcal{T}_1 and \mathcal{T}_2 are distinct, either $\mathcal{T}_1 - \mathcal{T}_2 \neq \emptyset$ or $\mathcal{T}_2 - \mathcal{T}_1 \neq \emptyset$; without loss of generality, there exists $X_1 \in \mathcal{T}_1 - \mathcal{T}_2$. Let $X_2 = S - X_1$. Then (X_1, X_2) is a $(\theta - 1)$ -separation, so one of X_1 or X_2 is in \mathcal{T}_2 by tangle axiom (T2). Thus, $X_2 \in \mathcal{T}_2$, so (X_1, X_2) is a $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation. □

We are particularly interested in $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separations, (X_1, X_2) , such that, for each $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation (Y_1, Y_2) , $\lambda(X_1) \leq \lambda(Y_1)$; we call (X_1, X_2) a *minimum-order $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation*.

3.4 Tree of tangles

Robertson and Seymour [25] (see also [11]) showed that a connectivity system with at least one tangle has a tree-decomposition displaying all of its *maximal* tangles (that is, tangles which are not contained in any other tangle).

Theorem 3.4.1. *If (S, λ) is a connectivity system, then (S, λ) has a tree-decomposition (T, μ) satisfying the following conditions:*

1. *for each maximal tangle \mathcal{T} in (S, λ) , there exists a unique node $t_{\mathcal{T}}$ in T displaying \mathcal{T} , and*
2. *for each non-leaf node t in T , there exists a unique maximal tangle \mathcal{T} in (S, λ) such that $t = t_{\mathcal{T}}$, and*
3. *for distinct maximal tangles \mathcal{T}_1 and \mathcal{T}_2 in (S, λ) , the path between $t_{\mathcal{T}_1}$ and $t_{\mathcal{T}_2}$ in T contains a tree-edge e presenting a minimum-order $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation.*

We prove here a similar result for tangles of some fixed order, θ . For $\theta \in \mathbb{N}$, if (S, λ) is a connectivity system, then a *tree of θ -tangles* in (S, λ) is a $(\theta - 1)$ -tree-decomposition (T, μ) in (S, λ) satisfying the following properties:

- (TT1)** for each node t in T of degree greater than three, there exists a unique θ -tangle \mathcal{T}_t in (S, λ) such that t displays \mathcal{T}_t ; and
- (TT2)** for distinct nodes t and t' in T each of degree greater than three, the path in T between t and t' contains an edge presenting a minimum-order $(\mathcal{T}_t, \mathcal{T}_{t'})$ -distinguishing separation.

A tree of θ -tangles, (T, μ) , in (S, λ) is a *full tree of θ -tangles* if it satisfies the following additional property:

- (FTT)** for each non-leaf node t in T , there exists a unique θ -tangle \mathcal{T} in (S, λ) such that $t = t_{\mathcal{T}}$.

Observe that property (FTT) implies property (TT1).

For $\theta \in \mathbb{N}$, a connectivity system (S, λ) is said to be *θ -elementary* if, for each $x \in S$, $\lambda(\{x\}) < \theta$.

The existence of a full tree of θ -tangles can be shown using Theorem 3.4.1.

Theorem 3.4.2. *For $\theta \in \mathbb{N}$, if (S, λ) is a θ -elementary connectivity system and (S, λ) contains a θ -tangle, then (S, λ) has a full tree of θ -tangles.*

Proof. For each θ -tangle \mathcal{T} , let $\tilde{\mathcal{T}}$ be a maximal tangle (of any order) containing \mathcal{T} . If \mathcal{T}_1 and \mathcal{T}_2 are distinct θ -tangles, then we saw in Lemma 3.3.1 that there exists a $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation (X_1, X_2) . Note that, by tangle property (T4'), $X_1 \in \tilde{\mathcal{T}}_1 - \tilde{\mathcal{T}}_2$, so $\tilde{\mathcal{T}}_1 \neq \tilde{\mathcal{T}}_2$. This also shows that (X_1, X_2) is a $(\tilde{\mathcal{T}}_1, \tilde{\mathcal{T}}_2)$ -distinguishing separation of order less than θ . Therefore, each minimum-order $(\tilde{\mathcal{T}}_1, \tilde{\mathcal{T}}_2)$ -distinguishing separation has order less than θ , and hence is a $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation.

Let $(\tilde{T}, \tilde{\mu})$ be the tree-decomposition shown to exist in Theorem 3.4.1. For each θ -tangle \mathcal{T} , let $\tilde{t}_{\mathcal{T}}$ be the node in \tilde{T} displaying the maximal tangle $\tilde{\mathcal{T}}$.

Let $E_{\theta} \subseteq E(\tilde{T})$ be a set of minimum cardinality such that, for each pair of distinct θ -tangles \mathcal{T}_1 and \mathcal{T}_2 , there exists a tree-edge $e_{\mathcal{T}_1, \mathcal{T}_2} \in E_{\theta}$ presenting a minimum-order $(\tilde{\mathcal{T}}_1, \tilde{\mathcal{T}}_2)$ -distinguishing separation; as we just showed, this is also a minimum-order $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation.

Let $E_{\ell} \subseteq E(\tilde{T})$ be the set of tree-edges in T incident with an edge. Let

$$T = \tilde{T} / (E(\tilde{T}) - (E_{\theta} \cup E_{\ell})).$$

Because no tree-edge incident with a leaf was contracted, $L(T) = L(\tilde{T})$, so let $\mu = \tilde{\mu}$. Each edge in T presents the same separation in (T, μ) and $(\tilde{T}, \tilde{\mu})$.

Each separation presented by a tree-edge in (T, μ) is either a minimum-order distinguishing separation between two θ -tangles, or is of the form $(\{x\}, S - \{x\})$ for some $x \in S$; in either case, the edge has width less than θ , so (T, μ) is a $(\theta - 1)$ -tree-decomposition.

For distinct θ -tangles \mathcal{T}_1 and \mathcal{T}_2 , T contains a tree-edge, $e_{\mathcal{T}_1, \mathcal{T}_2}$, presenting a minimum-order $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation. By Lemma 3.2.1, $e_{\mathcal{T}_1, \mathcal{T}_2}$ is on the path between $t_{\mathcal{T}_1}$ and $t_{\mathcal{T}_2}$ in T , so (TT2) holds.

Claim 3.4.2.1. *Property (FTT) holds for (T, μ) .*

Proof of Claim. Suppose for contradiction that there exists a non-leaf tree-node t that does not display any θ -tangle.

Because (S, λ) has at least one θ -tangle, \mathcal{T}_0 , and \mathcal{T}_0 is presented at some non-leaf node distinct from t , t is incident with at least one tree-edge e such that e is not incident with a leaf. Let k be the maximum width of a tree-edge incident with t and not incident with a leaf.

Note that, for each $x \in S$, T contains exactly one edge presenting the separation $(\{x\}, S - \{x\})$ —namely, the edge incident with $\mu(x)$. Therefore, for each tree-edge e not incident with a leaf, $e = e_{\mathcal{T}_1, \mathcal{T}_2}$, for some pair of θ -tangles, \mathcal{T}_1 and \mathcal{T}_2 .

Suppose \mathcal{T}_1 and \mathcal{T}_2 are distinct θ -tangles, $e_{\mathcal{T}_1, \mathcal{T}_2}$ is incident with t , and $e_{\mathcal{T}_1, \mathcal{T}_2}$ has width equal to k . Note that e lies on the path between $t_{\mathcal{T}_1}$ and $t_{\mathcal{T}_2}$. Without loss of generality, the component of $T \setminus \{e\}$ containing t also contains $t_{\mathcal{T}_1}$ but not $t_{\mathcal{T}_2}$.

There exists some tree-edge e_1 incident with t such that $t_{\mathcal{T}_1}$ and t lie in distinct components of $T \setminus \{e_1\}$, and $e_1 \neq e$. Note that the separation presented by e_1 is also a $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation, so e_1 has width no smaller than the width of e ; but e was chosen to have maximum width, so the width of e_1 equals the width of e , and, hence, the separation presented by e_1 is a minimum-order $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation. Thus, T contains two distinct edges each presenting a minimum-order $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation. This contradicts minimality of $|E_\theta|$. \square (Claim)

Property (TT1) follows from (FTT), so (T, μ) is a full tree of θ -tangles. \square

3.4.1 Bounded-degree tree of θ -tangles

For $\theta \in \mathbb{N}$, if (S, λ) is a connectivity system, then a *bounded-degree tree of θ -tangles* is a tree of θ -tangles, (T, μ) such that, for each node $t \in V(T)$ with $\deg_T(t) > 3$, t displays a θ -tangle \mathcal{T}_t with covering-number δ_t such that $\deg_T(t) \leq \theta\delta_t$.

We say that a connectivity system (S, λ) is *initial* if it satisfies the following property:

(INIT) for each separation (X, Y) in (S, λ) and each $x \in S$, if $|X| > 1$ and $|Y| > 1$, then $\lambda(X) > \lambda(\{x\})$.

Theorem 3.4.3. *For $\theta \in \mathbb{N}$, if (S, λ) is a θ -elementary, initial connectivity system, then (S, λ) has a bounded-degree tree of θ -tangles.*

Property (INIT), while slightly cumbersome, does not add any serious restriction for applications to graphs. We show now that Theorem 3.1.1, restated here, follows from Theorem 3.4.3.

Theorem 3.1.1. *For $\theta \in \mathbb{N}$ with $\theta > 2$, if G is a graph, then there exists a $(\theta - 1)$ -tree-decomposition (T, μ) of G such that for each node t in T with $\deg_T(t) > 3$, G contains a unique θ -tangle \mathcal{T}_t such that \mathcal{T}_t is displayed by t ; moreover, if δ_t is the covering number of \mathcal{T}_t , then*

$$\deg_T(t) \leq \theta\delta_t.$$

Proof. Suppose the contrary and let G be a minor-minimal graph with no bounded-degree tree of θ -tangles. If G is 3-connected then $(E(G), \lambda_G)$ satisfies property (INIT); because $\theta > 2$, $(E(G), \lambda_G)$ is θ -elementary, so, by Theorem 3.4.3, the desired tree of θ -tangles exists.

Otherwise, G is constructed from two proper minors G_1 and G_2 by a clique-sum of order at most two. By minor-minimality of G , G_1 and G_2 each have a bounded-degree tree of θ -tangles. These can be combined along the clique-sum to construct a bounded-degree tree of θ -tangles for G . \square

The factor of θ in the definition of a bounded-degree tree of θ -tangles is due to a technical limitation of the proof and is likely unnecessary; we conjecture that the following exact version of Theorem 3.4.3 is also true.

Conjecture 3.4.4. *For $\theta \in \mathbb{N}$, if (S, λ) is a θ -elementary, initial connectivity system, then (S, λ) has a tree of θ -tangles such that, for each tree-node $t \in V(T)$ with $\deg_T(t) > 3$, the degree of t equals the covering-number of the unique θ -tangle displayed at t .*

3.5 Partial tree-decomposition

Recall from Section 2.6 that, if (S, λ) is a connectivity system and $\mathcal{A} \subseteq 2^S$, then a *partial tree-decomposition over \mathcal{A}* is a pair (T, μ) , where T is a tree and $\mu : S \rightarrow L(T)$ such that, for each leaf $t \in L(T)$, there exists some $A \in \mathcal{A}$ such that

$$\{x \in S : \mu(x) = t\} \subseteq A.$$

The following lemma generalizes Lemma 3.2.1 and shows how each θ -tangle is displayed at a unique node in each partial $(\theta - 1)$ -tree-decomposition.

Lemma 3.5.1. *If*

- (1) (S, λ) is a connectivity system
- (2) \mathcal{T} is a θ -tangle in (S, λ) ,
- (3) $\mathcal{A} \subseteq 2^S$, and
- (4) (T, μ) is a partial $(\theta - 1)$ -tree-decomposition over \mathcal{A} ,

then there exists a unique node $t_{\mathcal{T}}$ in T such that, for each incidence $(t, e) \in I(T)$, (t, e) presents a set in \mathcal{T} if and only if t and $t_{\mathcal{T}}$ are in the same component of $T \setminus \{e\}$; moreover, if $\mathcal{A} \subseteq \mathcal{T}$, then $t_{\mathcal{T}}$ is not a leaf, and is the unique node in T displaying \mathcal{T} .

Proof. Let \vec{T} be the orientation of T such that, for each incidence (t, e) presenting a set in \mathcal{T} , e is oriented toward t ; this is a well-defined orientation because exactly one side of the $(\theta - 1)$ -separation presented by each tree-edge e is in \mathcal{T} . Because \vec{T} is acyclic, there exists a node $t_{\mathcal{T}}$ such that each tree-edge incident with $t_{\mathcal{T}}$ is oriented toward $t_{\mathcal{T}}$.

Claim 3.5.1.1. *If t is a node incident with tree-edges e_1 and e_2 in T , then at least one of e_1 or e_2 is oriented toward t .*

Proof of Claim. Suppose e_1 and e_2 are both oriented away from t in \vec{T} and X_1 and X_2 are the sets presented by the incidences (t, e_1) and (t, e_2) respectively. Then $S - X_1$ and $S - X_2$ are both in \mathcal{T} . However, X_1 and X_2 are disjoint so

$$(S - X_1) \cup (S - X_2) = S,$$

which contradicts tangle property (T3). □ (Claim)

If t is a node in T and P is the shortest path in T from t to $t_{\mathcal{T}}$ then P contains an edge oriented toward $t_{\mathcal{T}}$ (the edge in P incident with $t_{\mathcal{T}}$), so, by Claim 3.5.1.1, each edge of P is oriented toward $t_{\mathcal{T}}$. Therefore, the only tree-edge incident with t that is oriented away from t is the edge in P . Thus, for each tree-edge e incident with t , (t, e) displays a separation in \mathcal{T} if and only if t and $t_{\mathcal{T}}$ are in the same component of $T \setminus \{e\}$.

If $\mathcal{A} \subseteq \mathcal{T}$, then $t_{\mathcal{T}}$ is not a leaf, because each leaf-incidence presents the complement of a set in \mathcal{A} ; therefore, $t_{\mathcal{T}}$ displays \mathcal{T} , and no other node in T displays \mathcal{T} . □

3.6 Partial tree of θ -tangles

For each connectivity system (S, λ) and each $\mathcal{A} \subseteq S$, a *partial tree of θ -tangles over \mathcal{A}* is a partial $(\theta - 1)$ -tree-decomposition satisfying the properties of trees of tangles, (TT1) and (TT2), repeated here, and the additional property (PTT):

(TT1) for each node t in T of degree greater than three, there exists a unique θ -tangle \mathcal{T}_t in (S, λ) such that t displays \mathcal{T}_t ; and

(TT2) for distinct nodes t and t' in T each of degree greater than three, the path in T between t and t' contains an edge presenting a minimum-order $(\mathcal{T}_t, \mathcal{T}_{t'})$ -distinguishing separation.

(PTT) for each node t in T of degree greater than three, $\mathcal{T}_t \supseteq \mathcal{A}$.

It follows directly from the definition that a partial tree of θ -tangles over the family of singletons is equivalent tree of θ -tangles.

Lemma 3.6.1. *For each $\theta \in \mathbb{N}$ and each connectivity system (S, λ) , (T, μ) is a partial tree of θ -tangles over $\mathcal{A} = \{\{x\} : x \in S\}$ in (S, λ) , if and only if (T, μ) is a tree of θ -tangles in (S, λ) .*

3.7 Robust sets

For each $m \in \mathbb{N}$ and each connectivity system (S, λ) , a set $X \subseteq S$ is m -robust if there is no partition (X_1, \dots, X_m) of X such that, for each $i \in \{1, \dots, m\}$, $\lambda(X_i) < \lambda(X)$ and $\lambda(X - X_i) \leq \lambda(X)$. A 2-robust set is simply a set that cannot be partitioned into two sets of smaller connectivity; Robertson and Seymour [25] call 2-robust sets *robust*. If X is m -robust for each $m \in \mathbb{N}$, then X is ∞ -robust. The simplest example of an ∞ -robust set is a singleton, but it also turns out that minimum-order distinguishing separations between tangles are ∞ -robust.

Lemma 3.7.1. *For $\theta \in \mathbb{N}$, if $\mathcal{T}_1, \mathcal{T}_2$ are distinct θ -tangles in a connectivity system (S, λ) , and (X_1, X_2) is a minimum-order $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation, then X_1 and X_2 are both ∞ -robust.*

Proof. Suppose for the sake of contradiction that there exists a minimum-order $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation (X_1, X_2) such that X_1 is not ∞ -robust, so X_1 is not m -robust for some $m \in \mathbb{N}$; choose (X_1, X_2) and m to minimize m . Let $\{Y_1, \dots, Y_m\}$ be a partition of X_1 such that for each $j \in \{1, \dots, m\}$, $\lambda(Y_j) < \lambda(X_1)$ and $\lambda(X_1 - Y_j) \leq \lambda(X_1)$.

Note that $X_2 \neq S$ by (T4), so $X_1 \neq \emptyset$, so $m > 0$.

Because $\lambda(Y_j) < \lambda(X_1)$ and $Y_j \subseteq X_1$, $Y_j \in \mathcal{T}_1$ but $(Y_j, S - Y_j)$ is not a $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation, so $Y_j \in \mathcal{T}_2$ as well. Therefore, $X_1 - Y_j \notin \mathcal{T}_2$ or else $\{Y_j, X_1 - Y_j, X_2\}$ would cover S by 3 sets in \mathcal{T}_2 . But $X_1 - Y_j \in \mathcal{T}_1$, so by minimality of (X_1, X_2) , $\lambda(X_1 - Y_j) \geq \lambda(X_1)$, so $\lambda(X_1 - Y_j) = \lambda(X_1)$. Therefore, $(X_1 - Y_j, S - (X_1 - Y_j))$ is another minimum-order $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation.

For each $j \in \{1, \dots, m-1\}$, $\lambda(Y_j) < \lambda(X_1) = \lambda(X_1 - Y_m)$ and $\lambda((X_1 - Y_m) - Y_j) \leq \lambda(X_1 - Y_m)$ because

$$\begin{aligned} & \lambda((X_1 - Y_m) - Y_j) + \lambda(X_1) \\ &= \lambda((X_1 - Y_m) \cap (X_1 - Y_j)) + \lambda((X_1 - Y_m) \cup (X_1 - Y_j)) \\ &\leq \lambda(X_1 - Y_m) + \lambda(X_1 - Y_j) \\ &\leq \lambda(X_1 - Y_m) + \lambda(X_1) \end{aligned}$$

Hence, $X_1 - Y_m$ is a minimum-order $(\mathcal{T}_1, \mathcal{T}_2)$ -distinguishing separation and $X_1 - Y_m$ is not $(m-1)$ -robust, contradicting minimality of m . \square

The key property of ∞ -robust sets that we require here is that, given a collection of disjoint ∞ -robust sets in a tangle \mathcal{T} , \mathcal{T} contains a cover of S that is only slightly larger than the covering-number of \mathcal{T} and does not cross any of the ∞ -robust sets. This can be thought of as “identifying” each of the ∞ -robust sets with a single element, in which case we are showing that the covering number of a tangle cannot increase too much when identifying a collection of disjoint ∞ -robust sets.

Lemma 3.7.2. *For $\theta \in \mathbb{N}$, if (S, λ) is a connectivity system, \mathcal{A} is a collection of disjoint, ∞ -robust sets in (S, λ) , and $\mathcal{T} \supseteq \mathcal{A}$ is a θ -tangle in (S, λ) with covering-number δ , then S can be covered by $n \leq \theta\delta$ disjoint sets $\{Y_1, \dots, Y_n\} \subseteq \mathcal{T}$ such that, for each $A \in \mathcal{A}$, $A \subseteq Y_i$ for some $i \in \{1, \dots, n\}$.*

Lemma 3.7.2 is a special case of Lemma 3.7.3.

Lemma 3.7.3. *For $\theta \in \mathbb{N}$, if*

1. (S, λ) is a connectivity system,
2. \mathcal{A} is a collection of disjoint, ∞ -robust sets in (S, λ) ,
3. $\mathcal{T} \supseteq \mathcal{A}$ is a θ -tangle in (S, λ) with covering-number δ
4. $d \in \{0, \dots, \delta\}$, and
5. $\{X_1, \dots, X_\delta\}$ is a cover of S by sets in \mathcal{T} such that, for each $i \in \{d+1, \dots, \delta\}$ and each $A \in \mathcal{A}$, $X_i \cap A \in \{\emptyset, A\}$,

then S can be covered by disjoint sets $\{Y_1, \dots, Y_n\} \subseteq \mathcal{T}$ such that for each $A \in \mathcal{A}$, $A \subseteq Y_i$ for some i , and $n - \delta \leq d(\theta - 1)$.

Proof. Choose $\{Y_1, \dots, Y_n\} \subseteq \mathcal{T}$ covering S and such that

- (a) for each $i \in \{d+1, \dots, n\}$, and each $A \in \mathcal{A}$, $Y_i \cap A \in \{\emptyset, A\}$ and
- (b) $\sum_{i=1}^d \lambda(Y_i) + n - \delta \leq d(\theta - 1)$ and
- (c) $|\{A \in \mathcal{A} : \exists i \in \{1, \dots, n\}, A \cap Y_i \notin \{\emptyset, A\}\}|$ is minimized subject to (a) and (b), and
- (d) $\sum_{i=1}^n |Y_i|$ is minimized subject to (a), (b) and (c).

Note that $\{X_1, \dots, X_\delta\}$ is a covering satisfying (a) and (b), and both functions being minimized have non-negative integer value, so $\{Y_1, \dots, Y_n\}$ exists.

Claim 3.7.3.1. *The sets in $\{Y_1, \dots, Y_n\}$ are pairwise disjoint.*

Proof of Claim. If $i, j \in \{1, \dots, n\}$ are distinct and $Y_i \cap Y_j \neq \emptyset$, then, by submodularity, either $\lambda(Y_i - Y_j) \leq \lambda(Y_i)$ or $\lambda(Y_j - Y_i) \leq \lambda(Y_j)$, so without loss of generality, $\lambda(Y_i - Y_j) \leq \lambda(Y_i)$. $A \cap (Y_i - Y_j) = \emptyset$ if $A \cap Y_i = \emptyset$ or $A \subseteq Y_j$; otherwise $A \subseteq Y_i$ and $A \cap Y_j = \emptyset$, so $A \subseteq Y_i - Y_j$. In both cases, $A \cap (Y_i - Y_j) \in \{A, \emptyset\}$, so replacing Y_i by $Y_i - Y_j$ cannot make (a) or (b) false or increase the objective function in (c), but decreases the objective function in (d), contradicting minimality. \square (Claim)

If

$$\{A \in \mathcal{A} : \exists i \in \{1, \dots, n\}, A \cap Y_i \notin \{\emptyset, A\}\} = \emptyset,$$

then, for each $A \in \mathcal{A}$ there exists $i \in \{1, \dots, n\}$ such that $A \subseteq Y_i$. By (b),

$$n - \delta \leq d(\theta - 1) - \sum_{i=1}^d \lambda(Y_i) \leq d(\theta - 1).$$

Hence, $\{Y_1, \dots, Y_n\}$ satisfy all of the required properties, and the lemma holds.

Now suppose that there exists $A_0 \in \mathcal{A}$ and $i \in \{1, \dots, n\}$ such that $A_0 \cap Y_i \notin \{\emptyset, A_0\}$. Then, for each $i \in \{1, \dots, n\}$, $A_0 \not\subseteq Y_i$, so, for each $i \in \{d+1, \dots, n\}$, $Y_i \cap A_0 = \emptyset$. Let

$$I_{\cup} = \{i \in \{1, \dots, d\} : \lambda(A_0 \cup Y_i) \leq \lambda(Y_i)\}$$

and

$$I_{-} = \{i \in \{1, \dots, d\} : \lambda(Y_i - A_0) \leq \lambda(Y_i)\}$$

Claim 3.7.3.2. $I_{\cup} \cup I_{-} = \{1, \dots, d\}$

Proof of Claim. If $i \notin I_{\cup}$ then, by submodularity, $\lambda(A_0 \cap Y_i) < \lambda(A_0)$, so by 2-robustness of A_0 , $\lambda(A_0 - Y_i) \geq \lambda(A_0)$, so, by submodularity, $\lambda(Y_i - A_0) \leq \lambda(Y_i)$, so $i \in I_-$. \square (Claim)

For $i \in \{1, \dots, n\}$, define $Y'_i \in \mathcal{T}$ as follows. If $i \in \{d+1, \dots, n\}$, then $Y'_i = Y_i$; if $i \in I_{\cup}$ then $Y'_i = Y_i \cup A_0$; if $i \in I_- - I_{\cup}$ then $Y'_i = Y_i - A_0$.

Claim 3.7.3.3. For $i \in \{1, \dots, d\}$, $Y'_i \in \mathcal{T}$

Proof of Claim. If $i \in I_{\cup}$ then either $Y'_i = Y_i \cup A_0 \in \mathcal{T}$ or $S - (A_0 \cup Y_i) \in \mathcal{T}$; but the latter cannot occur, or else $\{A_0, Y_i, S - (A_0 \cup Y_i)\}$ would be a cover of S by 3 sets in \mathcal{T} .

If $i \in I_-$ then $Y'_i = Y_i - A_0 \subseteq Y_i \in \mathcal{T}$, so $Y_i - A_0 \in \mathcal{T}$. \square (Claim)

Claim 3.7.3.4. $I_{\cup} = \emptyset$.

Proof of Claim. Note that $\{Y'_1, \dots, Y'_n\}$ satisfies (a) because, for each $i \in \{d+1, \dots, n\}$, $Y'_i = Y_i$. Note also that $\{Y'_1, \dots, Y'_n\}$ satisfies (b) because $\sum_{i=1}^d \lambda(Y'_i) \leq \sum_{i=1}^d \lambda(Y_i)$. If $A \in \mathcal{A}$, $i \in \{1, \dots, d\}$ and $A \cap Y'_i \notin \{\emptyset, A\}$, then $A \neq A_0$ by the choice of $\{Y'_1, \dots, Y'_n\}$, so $A \cap A_0 = \emptyset$. Therefore, $A \cap (Y_i \cup A_0) = A \cap (Y_i - A_0) = A \cap Y_i$, so $A \cap Y'_i = A \cap Y_i$. Hence,

$$\begin{aligned} \{A \in \mathcal{A} : \exists i \in \{1, \dots, d\}, A \cap Y'_i \notin \{\emptyset, A\}\} \subsetneq \\ \{A \in \mathcal{A} : \exists i \in \{1, \dots, d\}, A \cap Y_i \notin \{\emptyset, A\}\}. \end{aligned}$$

Therefore, by minimality of (c), $\{Y'_1, \dots, Y'_n\}$ must not cover S . Note that, for each $i \in \{1, \dots, n\}$, $Y'_i \subseteq Y_i - A_0$, so $S - A_0 \subseteq \bigcup_{i=1}^n Y'_i$. If $i \in I_{\cup}$ then $A_0 \subseteq Y'_i$, in which case $\{Y'_1, \dots, Y'_n\}$ would cover S , so $I_{\cup} = \emptyset$. \square (Claim)

Thus, for each $i \in \{1, \dots, n\}$, $\lambda(Y_i \cup A_0) > \lambda(Y_i)$, so, by submodularity, $\lambda(Y_i \cap A_0) < \lambda(A_0)$. Because Y_1, \dots, Y_n are pairwise disjoint, $\{Y_i \cap A_0 : i \in \{1, \dots, n\}\}$ is a partition of A_0 . Therefore, by ∞ -robustness of A_0 , there is some $i \in \{1, \dots, n\}$ such that $\lambda(A_0 - Y_i) > \lambda(A_0)$, so, by submodularity, $\lambda(Y_i - A_0) < \lambda(Y_i)$. Hence, $\sum_{i=1}^d \lambda(Y'_i) < \sum_{i=1}^d \lambda(Y_i)$.

Let $Y'_{n+1} = A_0$. Then $\{Y'_1, \dots, Y'_{n+1}\}$ is a cover of S by sets in \mathcal{T} . Property (a) is satisfied because $Y'_{n+1} = A_0$ is disjoint from the other sets in \mathcal{A} . Property (b) is satisfied because

$$\sum_{i=1}^d \lambda(Y'_i) + (n+1) - \delta \leq \sum_{i=1}^d \lambda(Y_i) + n - \delta \leq d(\theta - 1).$$

But this contradicts minimality of (c). \square

3.8 Tiebreakers

A *tiebreaker* for a connectivity system (S, λ) is a symmetric, submodular function $\hat{\lambda} : 2^S \rightarrow \mathbb{N}$ satisfying the following properties:

(TB1) for $X, Y \subseteq S$, if $\lambda(X) < \lambda(Y)$, then $\hat{\lambda}(X) < \hat{\lambda}(Y)$, and

(TB2) for distinct $X, Y \subseteq S$, $\hat{\lambda}(X) \neq \hat{\lambda}(Y)$.

Robertson and Seymour [25] proved that each connectivity system has a tiebreaker:

Lemma 3.8.1. *If (S, λ) is a connectivity system, then there exists a tiebreaker, $\hat{\lambda}$ for (S, λ) .*

If \mathcal{T} and \mathcal{T}' are distinct θ -tangles in a connectivity system (S, λ) with a tiebreaker $\hat{\lambda}$, then a $\hat{\lambda}$ -minimum $(\mathcal{T}, \mathcal{T}')$ -distinguishing separation is a $(\mathcal{T}, \mathcal{T}')$ -distinguishing separation (X, X') minimizing $\hat{\lambda}(X)$.

Lemma 3.8.2. *For $\theta \in \mathbb{N}$, if \mathcal{T} and \mathcal{T}' are distinct θ -tangles in a connectivity system (S, λ) with a tiebreaker $\hat{\lambda}$, and (X, X') is a $\hat{\lambda}$ -minimum $(\mathcal{T}, \mathcal{T}')$ -distinguishing separation, then for each partition (X_1, X_2) of X , with $X_1 \neq \emptyset$ and $X_2 \neq \emptyset$,*

$$\max\{\hat{\lambda}(X_1), \hat{\lambda}(X_2)\} > \hat{\lambda}(X).$$

Proof. Suppose the contrary, so, by (TB2), for $i \in \{1, 2\}$,

$$\hat{\lambda}(X_i) < \hat{\lambda}(X) < \theta;$$

Also

$$X_i \subseteq X \in \mathcal{T},$$

so $X_i \in \mathcal{T}$ by tangle property (T3). Because (X, X') is $\hat{\lambda}$ -minimum, $(X_i, S - X_i)$ is not a $(\mathcal{T}, \mathcal{T}')$ -distinguishing separation, so $X_i \in \mathcal{T}'$. Thus, $\{X_1, X_2, X'\}$ is a partition of S into three sets in \mathcal{T}' , contradicting tangle property (T4). \square

Robertson and Seymour [25] proved that tiebreakers give rise to “non-crossing” separations, in the following sense:

Lemma 3.8.3. *For $\theta \in \mathbb{N}$, if (S, λ) is a connectivity system with a tiebreaker $\hat{\lambda}$, $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}'_1, \mathcal{T}'_2$ are θ -tangles in (S, λ) , and (A_1, A_2) and (A'_1, A'_2) are $\hat{\lambda}$ -minimum $(\mathcal{T}_1, \mathcal{T}_2)$ - and $(\mathcal{T}'_1, \mathcal{T}'_2)$ -distinguishing separations, respectively, then, for some $i, j \in \{1, 2\}$, $A_i \cap A'_j = \emptyset$.*

In particular, we are interested in the following special case.

Lemma 3.8.4. *For $\theta \in \mathbb{N}$, if (S, λ) is a connectivity system with a tiebreaker $\hat{\lambda}$, $\mathcal{T}, \mathcal{T}_1, \mathcal{T}_2$ are θ -tangles in (S, λ) , and (A_1, B_1) and (A_2, B_2) are $\hat{\lambda}$ -minimum $(\mathcal{T}, \mathcal{T}_1)$ - and $(\mathcal{T}, \mathcal{T}_2)$ -distinguishing separations, respectively, then, either $A_1 \subseteq A_2$, or $A_2 \subseteq A_1$, or $A_1 \cap A_2 = \emptyset$.*

Proof. By Lemma 3.8.3 with $\mathcal{T}_1 = \mathcal{T}'_1 = \mathcal{T}$, one of the following occurs: $A_1 \cap A_2 = \emptyset$; $A_1 \cap B_2 = \emptyset$, in which case $A_1 \subseteq A_2$; $A_2 \cap B_1 = \emptyset$, in which case $A_2 \subseteq A_1$; $B_1 \cap B_2 = \emptyset$, in which case $A_1 \cup A_2 = S$, which contradicts tangle property (T4'). \square

3.9 Distinguishing covering

For $\theta \in \mathbb{N}$, if (S, λ) is a connectivity system with a tiebreaker $\hat{\lambda}$ and \mathcal{T} is a θ -tangle in (S, λ) , then a $\hat{\lambda}$ -distinguishing covering of \mathcal{T} is a family $\mathcal{Y} \subseteq \mathcal{T}$ satisfying the following properties:

(DC1) $\bigcup_{Y \in \mathcal{Y}} Y = S$;

(DC2) for each θ -tangle \mathcal{T}' in (S, λ) with $\mathcal{T}' \neq \mathcal{T}$, there exists a $\hat{\lambda}$ -minimum $(\mathcal{T}, \mathcal{T}')$ -distinguishing separation (X, X') such that, for each $Y \in \mathcal{Y}$, either $X \subseteq Y$ or $X \cap Y = \emptyset$.

Using Lemma 3.7.2, the existence of a distinguishing covering with small cardinality can be established.

Lemma 3.9.1. *For $\theta \in \mathbb{N}$, if (S, λ) is a connectivity system with a tiebreaker $\hat{\lambda}$ and \mathcal{T} is a θ -tangle in (S, λ) with covering-number δ , then there exists a $\hat{\lambda}$ -distinguishing covering \mathcal{Y} for \mathcal{T} such that $|\mathcal{Y}| \leq \theta\delta$.*

Proof. For each θ -tangle \mathcal{T}' distinct from \mathcal{T} , let $(A_{\mathcal{T}'}, B_{\mathcal{T}'})$ be a $\hat{\lambda}$ -minimum $(\mathcal{T}, \mathcal{T}')$ -distinguishing separation. Let \mathcal{A}_0 be the family consisting of $A_{\mathcal{T}'}$ for each θ -tangle \mathcal{T}' distinct from \mathcal{T} . Let $\mathcal{A}_1 = \{\{x\} : x \in S\}$. Let \mathcal{A} be the maximal sets in $\mathcal{A}_0 \cup \mathcal{A}_1$. By Lemma 3.8.4, no pair of sets in \mathcal{A} can properly intersect, so \mathcal{A} is a partition of S .

By Lemma 3.7.1, \mathcal{A} is a family of ∞ -robust sets. By, Lemma 3.7.2, there exists a partition \mathcal{Y} of S such that $|\mathcal{Y}| \leq \theta\delta$, $\mathcal{Y} \subseteq \mathcal{T}$ and, for each $A \in \mathcal{A}$, $A \subseteq Y$ for some $Y \in \mathcal{Y}$.

Therefore, for each θ -tangle \mathcal{T}' distinct from \mathcal{T} , there is some $A \in \mathcal{A}$ and some $Y \in \mathcal{Y}$ such that

$$A_{\mathcal{T}'} \subseteq A' \subseteq Y;$$

moreover, because Y is a partition, for $Y' \in \mathcal{Y} - \{Y\}$, $A_{\mathcal{T}'} \cap Y' = \emptyset$. Thus, \mathcal{Y} is a $\hat{\lambda}$ -distinguishing covering, as desired. \square

If $\hat{\lambda}$ is a tiebreaker for (S, λ) , then a $\hat{\lambda}$ -distinguishing covering \mathcal{Y} for \mathcal{T} is said to be *minimum* if it satisfies the following property:

(MDC) for each distinguishing covering \mathcal{Y}' for \mathcal{T} with $|\mathcal{Y}'| \leq |\mathcal{Y}|$,

$$\sum_{Y \in \mathcal{Y}} \hat{\lambda}(Y) \leq \sum_{Y' \in \mathcal{Y}'} \hat{\lambda}(Y')$$

Next we show some important properties of minimum distinguishing coverings.

Lemma 3.9.2. *For $\theta \in \mathbb{N}$, if \mathcal{Y} is a $\hat{\lambda}$ -minimum distinguishing covering for a θ -tangle \mathcal{T} in a connectivity system (S, λ) with a tiebreaker $\hat{\lambda}$, then \mathcal{Y} is a partition of S .*

Proof. By (DC1), it suffices to show that the sets in \mathcal{Y} are pairwise disjoint. Suppose $Y_1, Y_2 \in \mathcal{Y}$ are distinct and $Y_1 \cap Y_2 \neq \emptyset$. Because $\hat{\lambda}$ is a tiebreaker, we may assume $\hat{\lambda}(Y_1 - Y_2) < \hat{\lambda}(Y_1)$. Let $\mathcal{Y}' = (\mathcal{Y} - \{Y_1\}) \cup \{Y_1 - Y_2\}$.

Claim 3.9.2.1. *The family \mathcal{Y}' is a distinguishing covering for \mathcal{T} .*

Proof of Claim. Note first that

$$\lambda(Y_1 - Y_2) \leq \lambda(Y_1) < \theta,$$

and $Y_1 - Y_2 \subseteq Y_1$, so $Y_1 - Y_2 \in \mathcal{T}$ by property (T3). Note also that $\bigcup_{Y \in \mathcal{Y}'} Y = S$, so (DC1) holds. To prove (DC2), suppose that \mathcal{T}' is a θ -tangle in (S, λ) with $\mathcal{T}' \neq \mathcal{T}$. By property (DC2) of \mathcal{Y} , there exists a minimum-order $(\mathcal{T}, \mathcal{T}')$ -distinguishing separation (X, X') such that, for each $Y \in \mathcal{Y}$, either $X \subseteq Y$ or $X \cap Y = \emptyset$. If $X \subseteq Y_1$ and $X \subseteq Y_2$ then $X \subseteq Y_1 \cap Y_2$, so $X \cap (Y_1 - Y_2) = \emptyset$. Otherwise, $X \cap (Y_1 \cap Y_2) = \emptyset$ so $X \cap (Y_1 - Y_2) = X \cap Y_1$. In either case, (DC2) holds, and \mathcal{Y}' is a distinguishing covering for \mathcal{T} . \square (Claim)

Also, we have that

$$\sum_{Y' \in \mathcal{Y}'} \hat{\lambda}(Y') < \sum_{Y \in \mathcal{Y}} \hat{\lambda}(Y),$$

contradicting (MDC). \square

Lemma 3.9.3. For $\theta \in \mathbb{N}$, if

1. (S, λ) is a connectivity system with a tiebreaker $\hat{\lambda}$,
2. \mathcal{T} is a θ -tangle in (S, λ) ,
3. \mathcal{Y} is a $\hat{\lambda}$ -minimum distinguishing covering for (S, λ) ,
4. $Y \in \mathcal{Y}$,
5. $Z \in \mathcal{T}$ such that $Y \subseteq Z$,

then $\hat{\lambda}(Z) \geq \hat{\lambda}(Y)$.

Proof. For each θ -tangle \mathcal{T}' in (S, λ) with $\mathcal{T}' \neq \mathcal{T}$, let $X_{\mathcal{T}'} \in \mathcal{T}$ such that $(X_{\mathcal{T}'}, S - X_{\mathcal{T}'})$ is a $\hat{\lambda}$ -minimum $(\mathcal{T}, \mathcal{T}')$ -distinguishing separation and for each $Y \in \mathcal{Y}$, either $X_{\mathcal{T}'} \subseteq Y$ or $X_{\mathcal{T}'} \cap Y = \emptyset$.

Suppose there exists $Z \in \mathcal{T}$ such that $Y \subseteq Z$ and $\hat{\lambda}(Z) < \hat{\lambda}(Y)$; choose such an Z minimizing $\hat{\lambda}(Z)$.

If, for each θ -tangle \mathcal{T}' , either $X_{\mathcal{T}'} \subseteq Z$ or $X_{\mathcal{T}'} \cap Z = \emptyset$, then replacing Y by Z yields a distinguishing covering that contradicts minimality in (MDC). Therefore, we may assume that there exists a θ -tangle \mathcal{T}' such that $X_{\mathcal{T}'} \cap Z \notin \{\emptyset, X_{\mathcal{T}'}\}$.

By minimality of $\hat{\lambda}(Z)$, $\hat{\lambda}(Z \cup X_{\mathcal{T}'}) > \hat{\lambda}(Z)$.

By submodularity, $\hat{\lambda}(Z \cap X_{\mathcal{T}'}) < \hat{\lambda}(X_{\mathcal{T}'})$. By Lemma 3.8.2 applied to $X_{\mathcal{T}'}$, $\hat{\lambda}(X_{\mathcal{T}'} - Z) \geq \hat{\lambda}(X_{\mathcal{T}'})$; because $\hat{\lambda}$ is a tiebreaker, $\hat{\lambda}(X_{\mathcal{T}'} - Z) > \hat{\lambda}(X_{\mathcal{T}'})$. By submodularity, $\hat{\lambda}(Z - X_{\mathcal{T}'}) < \hat{\lambda}(Z)$. However, $X_{\mathcal{T}'} \not\subseteq Y$ because $X_{\mathcal{T}'} \not\subseteq Z$ and $Y \subseteq Z$, so $X_{\mathcal{T}'} \cap Y = \emptyset$. Therefore, $Y \subseteq Z - X_{\mathcal{T}'}$, contradicting the choice of Z . \square

3.10 Constructing a bounded-degree tree of θ -tangles

For $\theta \in \mathbb{N}$, if (S, λ) is a connectivity system and $\mathcal{A} \subseteq 2^S$, then a partial tree of θ -tangles (T, μ) over \mathcal{A} is a *bounded-degree partial tree of θ -tangles over \mathcal{A} in (S, λ)* if and only if (T, μ) satisfies the following property:

(BDPTT) for each node $t \in T$, if $\deg_T(t) > 3$ and δ_t is the covering-number of the unique θ -tangle displayed by t , then $\deg_T(t) \leq \theta \delta_t$.

We say that a partial $(\theta - 1)$ -tree-decomposition, (T, μ) over \mathcal{A} is *tight* if it satisfies the following property:

(TIGHT) for each leaf $t \in L(T)$,

$$\{x \in S : \mu(x) = t\} \in \mathcal{A}.$$

If (S, λ) is a connectivity system, then a family $\mathcal{A} \subseteq 2^S$ is said to be a $\hat{\lambda}$ -*monotone covering* of (S, λ) if it satisfies the following properties:

(MC1) $\bigcup_{A \in \mathcal{A}} A = S$;

(MC2) for $A, A' \in \mathcal{A}$, if $\hat{\lambda}(A) < \hat{\lambda}(A')$, then $A' \not\subseteq A$ and $S - A' \not\subseteq A$; and

(MC3) for $X \subseteq S$, if $\hat{\lambda}(X) \leq \max\{\hat{\lambda}(A) : A \in \mathcal{A}\}$, then there exists $A \in \mathcal{A}$ such that $\hat{\lambda}(A) \leq \hat{\lambda}(X)$ and either $X \subseteq A$ or $S - X \subseteq A$,

The *width* of a $\hat{\lambda}$ -monotone covering \mathcal{A} is $\max\{\lambda(A) : A \in \mathcal{A}\}$. A monotone covering \mathcal{A} is said to be θ -*complete* if it has width less than θ and satisfies the following property

(CMC) for $X \subseteq S$, if $\hat{\lambda}(X) < \theta$, then there exists $A \in \mathcal{A}$ such that $\hat{\lambda}(A) \leq \hat{\lambda}(X)$ and either $X \subseteq A$ or $S - X \subseteq A$.

Lemma 3.10.1. *For $\theta \in \mathbb{N}$, if \mathcal{A} is a θ -complete $\hat{\lambda}$ -monotone covering for a connectivity system (S, λ) and \mathcal{A} does not extend to any θ -tangle in (S, λ) , then (S, λ) has a tight, bounded-degree partial tree of θ -tangles over \mathcal{A} .*

Proof. Let \mathcal{T} be the family consisting of the sets $X \subseteq S$ such that $\lambda(X) < \theta$ and there exists $A \in \mathcal{A}$ with $X \subseteq A$. Let $\mathcal{Y} \subseteq \mathcal{T}$ such that

(i) $\bigcup_{Y \in \mathcal{Y}} Y = S$,

(ii) $|\mathcal{Y}|$ is minimum subject to (i),

(iii) $\sum_{Y \in \mathcal{Y}} \hat{\lambda}(Y_i)$ is minimum subject to (ii).

Claim 3.10.1.1. $|\mathcal{Y}| \leq 3$.

Proof of Claim. The family \mathcal{T} satisfies tangle properties (T1), by construction, (T2), by (CMC), and (T3), by construction.

If \mathcal{T} does not satisfy tangle property (T5), then there exists $x \in S$ such that $S - \{x\} \in \mathcal{T}$; by (MC1), there exists $A \in \mathcal{A}$ such that $x \in A$; therefore $\{S - \{x\}, A\} \subseteq \mathcal{T}$ and $S - \{x\} \cup A = S$, so $|\mathcal{Y}| \leq 2$.

Otherwise, because \mathcal{T} is not a θ -tangle property (T4) fails, so there is a partition $\{Y_1, Y_2, Y_3\}$ of S such that $\{Y_1, Y_2, Y_3\} \subseteq \mathcal{T}$; hence, $|\mathcal{Y}| \leq 3$. \square (Claim)

Claim 3.10.1.2. \mathcal{Y} is a partition of S .

Proof of Claim. If $Y \cap Y' \neq \emptyset$ for distinct Y and Y' in \mathcal{Y} , then, by submodularity, without loss of generality, $\hat{\lambda}(Y - Y') < \hat{\lambda}(Y)$. Therefore, replacing Y by $Y - Y'$ preserves (i) and (ii) but decreases the objective in (iii), contradicting the choice of \mathcal{Y} . \square (Claim)

Claim 3.10.1.3. $\mathcal{Y} \subseteq \mathcal{A}$.

Proof of Claim. Let $Y \in \mathcal{Y}$. By the definition of \mathcal{T} , there exists $A' \in \mathcal{T}$ such that $Y \subseteq A'$. If $Y = A'$ then $Y \in \mathcal{A}$, as required, so we may assume $Y \subsetneq A'$. Therefore, by (iii), $\hat{\lambda}(Y) < \hat{\lambda}(A')$.

By (CMC) there exists $A \in \mathcal{A}$ such that $\hat{\lambda}(A) \leq \hat{\lambda}(Y)$ and either $Y \subseteq A$ or $S - Y \subseteq A$. In the former case, by (iii), $Y = A$, so we may assume $S - Y \subseteq A$. But

$$\hat{\lambda}(A') > \hat{\lambda}(Y) \geq \hat{\lambda}(A),$$

and

$$S - A' \subseteq S - Y \subseteq A,$$

contradicting (MC2). \square (Claim)

Let T be a degree- $|\mathcal{Y}|$ star with set of leaves $\{t_Y : Y \in \mathcal{Y}\}$. For $Y \in \mathcal{Y}$ and $x \in Y$, let $\mu(x) = t_Y$. Then (T, μ) is a tight partial tight $(\theta - 1)$ -tree-decomposition over \mathcal{A} . Moreover, because T has no node of degree greater than three, (T, μ) is a tight, bounded-degree partial tree of θ -tangles over \mathcal{A} . \square

Lemma 3.10.2. For $\theta \in \mathbb{N}$, if \mathcal{A} is a θ -complete $\hat{\lambda}$ -monotone covering for a connectivity system (S, λ) and \mathcal{T} is a θ -tangle in (S, λ) such that $\mathcal{A} \subseteq \mathcal{T}$, then (S, λ) has a tight, bounded-degree partial tree of θ -tangles over \mathcal{A} .

Proof. Let δ be the covering-number of \mathcal{T} . By Lemma 3.9.1, \mathcal{T} has a distinguishing covering \mathcal{Y} with $|\mathcal{Y}| \leq \theta\delta$; we may assume that \mathcal{Y} is a minimum distinguishing covering. By Lemma 3.9.2, \mathcal{Y} is a partition.

Claim 3.10.2.1. $\mathcal{Y} \subseteq \mathcal{A}$.

Proof of Claim. For $Y \in \mathcal{Y}$, by (CMC), there exists $A \in \mathcal{A}$ such that $\hat{\lambda}(A) \leq \hat{\lambda}(Y)$ and either $Y \subseteq A$ or $S - Y \subseteq A$; the latter is impossible because both A and Y are in \mathcal{T} , so $A \cup Y \neq S$. Therefore $Y \subseteq A$. By Lemma 3.9.3, $\hat{\lambda}(Y) \leq \hat{\lambda}(A)$, so $\hat{\lambda}(Y) = \hat{\lambda}(A)$. Therefore, because $\hat{\lambda}$ is a tiebreaker, $Y = A$. \square (Claim)

Claim 3.10.2.2. For each θ -tangle \mathcal{T}' distinct from \mathcal{T} , $\mathcal{Y} \not\subseteq \mathcal{T}'$.

Proof of Claim. By (DC1) and (DC2), there exists a $\hat{\lambda}$ -minimum $(\mathcal{T}, \mathcal{T}')$ -distinguishing separation (X, X') and $Y \in \mathcal{Y}$ such that $X \subseteq Y$. Then

$$S - Y \subseteq X' \in \mathcal{T}',$$

so $S - Y \in \mathcal{T}'$ by tangle property (T3), so, by tangle property (T4), $Y \notin \mathcal{T}'$. \square (Claim)

Let T be a degree- $|\mathcal{Y}|$ star with set of leaves $\{t_Y : Y \in \mathcal{Y}\}$ adjacent to a degree- $|\mathcal{Y}|$ node t_0 . For $Y \in \mathcal{Y}$ and $y \in Y$, let $\mu(y) = t_Y$. Then (T, μ) is a tight partial $(\theta - 1)$ -tree-decomposition over \mathcal{A} . The node t_0 has degree at most $\theta\delta$ displays \mathcal{T} and no other θ -tangle, so (T, μ) is a tight, bounded-degree partial tree of θ -tangles over \mathcal{A} . \square

Lemma 3.10.3. For $\theta \in \mathbb{N}$, if \mathcal{A} is a $\hat{\lambda}$ -monotone covering of width less than θ for a connectivity system (S, λ) , then (S, λ) has a tight, bounded-degree partial tree of θ -tangles over \mathcal{A} .

Proof. Suppose the contrary and let $\mathcal{A} \subseteq 2^S$ be maximal such that \mathcal{A} is $\hat{\lambda}$ -monotone of width less than θ and (S, λ) has no tight, bounded-degree partial tree of θ -tangles.

Claim 3.10.3.1. If, for each $(\theta - 1)$ -separation (X_1, X_2) in S , there exists an $A \in \mathcal{A}$ such that either $X_1 \subseteq A$ or $X_2 \subseteq A$, then property (CMC) holds.

Proof of Claim. Let (X_1, X_2) be a $(\theta - 1)$ -separation. By assumption, there exists $A \in \mathcal{A}$ such that either $X_1 \subseteq A$ or $X_2 \subseteq A$. If $\hat{\lambda}(A) \leq \hat{\lambda}(X_1)$, then the claim holds, so we may assume $\hat{\lambda}(A) > \hat{\lambda}(X_1)$. Therefore,

$$\hat{\lambda}(X_1) \leq \max\{\hat{\lambda}(A) : A \in \mathcal{A}\},$$

so, by (MC3), there exists $A' \in \mathcal{A}$ such that $\hat{\lambda}(A') \leq \hat{\lambda}(X_1)$ and either $X_1 \subseteq A'$ or $X_2 \subseteq A'$, as desired. \square (Claim)

Claim 3.10.3.2. *There exists a $(\theta - 1)$ -separation (X_1, X_2) such that for each $A \in \mathcal{A}$, $X_1 \not\subseteq A$ and $X_2 \not\subseteq A$.*

Proof of Claim. By Claim 3.10.3.1, if there is no $(\theta - 1)$ -separation (X_1, X_2) such that, for $A \in \mathcal{A}$, $X_1 \not\subseteq A$ and $X_2 \not\subseteq A$, then the lemma follows from Lemma 3.10.1, in the case when \mathcal{A} does not extend to a θ -tangle, and Lemma 3.10.2, in the case when \mathcal{A} does extend to a θ -tangle. \square (Claim)

Choose a $(\theta - 1)$ -separation (X_1, X_2) such that, for each $A \in \mathcal{A}$, $X_1 \not\subseteq A$ and $X_2 \not\subseteq A$, and minimizing $\hat{\lambda}(X_1)$.

Claim 3.10.3.3. *For $i \in \{1, 2\}$, $\mathcal{A} \cup \{X_i\}$ is $\hat{\lambda}$ -monotone of width less than θ .*

Proof of Claim. Property (MC1) is trivial.

To prove property (MC2), note that, by (MC3) for \mathcal{A} and the choice of (X_1, X_2) , for each $A \in \mathcal{A}$, $\hat{\lambda}(X) > \hat{\lambda}(A)$, and, for each $A \in \mathcal{A}$, $X_1 \not\subseteq A$ and $X_2 \not\subseteq A$.

To prove property (MC3), let $X' \subseteq S$ such that $\hat{\lambda}(X') \leq \max\{\hat{\lambda}(A) : A \in \mathcal{A} \cup \{X_i\}\}$. If $\hat{\lambda}(X') < \max\{\hat{\lambda}(A) : A \in \mathcal{A}\}$, then, by (MC3), there exists $A \in \mathcal{A}$ such that $\hat{\lambda}(A) \leq \hat{\lambda}(X')$ and either $X' \subseteq A$ or $S - X' \subseteq A$. Otherwise, $\hat{\lambda}(X') \leq \hat{\lambda}(X_i)$, so, by the choice of X_i , either $X' = X_i$ or there exists $A \in \mathcal{A}$ such that $\hat{\lambda}(A) \leq \hat{\lambda}(X')$ and either $X' \subseteq A$ or $S - X' \subseteq A$. In each case, $\mathcal{A} \cup \{X_i\}$ satisfies property (MC3).

The width is less than θ because X_i is a $(\theta - 1)$ -separation. \square (Claim)

By the choice of \mathcal{A} , there exist tight, bounded-degree partial trees of θ -tangles, (T_1, μ_1) and (T_2, μ_2) over $\mathcal{A} \cup \{X_1\}$ and $\mathcal{A} \cup \{X_2\}$, respectively.

For $i \in \{1, 2\}$, if (T_i, μ_i) is a partial tree-decomposition over \mathcal{A} , then (T_i, μ_i) , trivially satisfies properties (TIGHT), (BDPTT), (PTT), (TT1) and (TT2), so (T_i, μ_i) is a tight, bounded-degree partial tree of θ -tangles over \mathcal{A} , and the lemma holds. Therefore, we may assume that (T_i, μ_i) has some leaf, t_i such that

$$\{\{x\} : \mu_i(x) = t_i\} = X_i.$$

Let T be obtained from the disjoint union of T_1 and T_2 by identifying vertices t_1 and t_2 . Let $\mu : S \rightarrow L(T)$ be defined as

$$\mu(x) = \begin{cases} \mu_1(x) & x \in X_2 \\ \mu_2(x) & x \in X_1 \end{cases}$$

Then (T, μ) is a tight, bounded-degree partial tree of θ -tangles over \mathcal{A} , as desired. \square

Theorem 3.4.3, restated here, now follows easily.

Theorem 3.4.3. *For $\theta \in \mathbb{N}$, if (S, λ) is a θ -elementary, initial connectivity system, then (S, λ) has a bounded-degree tree of θ -tangles.*

Proof. Let $\hat{\lambda}$ be a tiebreaker for (S, λ) . Let $\mathcal{A} = \{\{x\} : x \in S\}$. Then \mathcal{A} trivially satisfies properties (MC1) and (MC2) and satisfies property (MC3) by (INIT). Therefore, \mathcal{A} is a $\hat{\lambda}$ -monotone covering. By Lemma 3.10.3, there exists a tight, bounded-degree partial tree of θ -tangles, (T, μ) , over \mathcal{A} . Then (T, μ) is a bounded-degree tree of θ -tangles. \square

Chapter 4

Unavoidable-minors

4.1 Introduction

In this chapter we describe the “unavoidable-minors” arising from “large θ -connected pieces” of a graph. We begin with some examples.

4.1.1 Unavoidable-minors

It is easy to show, and has been previously observed in the literature [21], that each sufficiently large, connected graph contains either a high degree vertex or a long path. Hence, every sufficiently large connected graph has a minor that is either a high-degree star or a long path; see Figure 4.1.

Theorem 4.1.1. *For each $n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that if G is a connected graph with at least N vertices, then G contains a minor isomorphic to either a degree- n star or a path of length n .*

For each $n \in \mathbb{N}$, it is easy to show that each sufficiently large 2-connected graph contains either a cycle of length n or a $K_{2,n}$ -minor; see Figure 4.2. This has also been observed previously [21].

Theorem 4.1.2. *For each $n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that if G is a 2-connected graph with at least N vertices, then G contains either a $K_{2,n}$ -minor or a cycle with at least n edges.*




Figure 4.1: The unavoidable-minors in large connected graphs




Figure 4.2: The unavoidable-minors in large 2-connected graphs




Figure 4.3: The unavoidable-minors in large 3-connected graphs, depicted with $n = 12$


Oporowski, Oxley and Thomas [21] described the minors that are forced to appear in sufficiently large 3-connected graphs; a *wheel* of length n is constructed from a cycle of length n by adding a new vertex adjacent to every vertex in the cycle; see Figure 4.3.

Theorem 4.1.3. *For $n \in \mathbb{N}$, there exists $N \in \mathbb{N}$ such that, if G is a 3-connected graph with at least N vertices, then G contains either $K_{3,n}$ -minor or a length- n -wheel-minor.*


Oporowski, Oxley and Thomas [21] also described a family of minors that are forced to appear in each sufficiently large 4-connected graph. For each $n > 2$, a *2-wheel of length n* is a graph obtained from a cycle of length n by adding two new vertices that are adjacent to every vertex in the cycle, but are not adjacent to each other. A *cyclic-ladder of length n* is a graph obtained from two disjoint cycles of length n , with vertices u_1, \dots, u_n and v_1, \dots, v_n respectively appearing in that order around the cycles, by adding an edge between u_i and v_i for each $i \in \{1, \dots, n\}$. A *Möbius-ladder of length n* is a graph obtained from a cyclic ladder of length n by deleting the edges $\{u_1u_n, v_1v_n\}$ and adding the edges $\{u_1v_n, v_1, u_n\}$. See Figure 4.4.

Theorem 4.1.4. *For each $n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that if G is a 4-connected graph with at least N vertices, then G contains a minor isomorphic to one of the following: $K_{4,n}$, a 2-wheel of length n , a cyclic zig-zag ladder of length n , or a Möbius zig-zag ladder of length n .*


In this chapter we generalize Theorems 4.1.1, 4.1.2, 4.1.3 and 4.1.4.




(a) $K_{4,n}$



(b) A 2-wheel of length n



(c) A cyclic ladder of length n



(d) A Möbius ladder of length n

Figure 4.4: Examples of the unavoidable-minors for internally-4-connected graphs, depicted with $n = 12$

4.1.2 θ -connected sets

For $\theta \in \mathbb{N}$, a set of vertices X in a graph G is said to be a θ -connected set if, for each pair of subsets $Y, Z \subseteq X$ with $|Y| = |Z| \leq \theta$, G contains a collection of $|Y|$ vertex-disjoint paths between Y and Z .

Robertson, Seymour and Thomas [28] showed that the existence of a large, highly-connected sets of vertices in a graph is qualitatively equivalent to the existence of a high-order tangle; they defined the *tangle matroid* and showed that the independent sets in this matroid are fully connected to each other, which implies that the independent sets are highly connected.

Diestel, Jensen, Gorbonov and Thomassen [4] showed the following more direct connection between large, highly-connected sets of vertices and high tree-width.

Theorem 4.1.5. *For each natural number ω , and for each graph G ,*

- (i) *if G contains an $(\omega + 1)$ -connected set of size at least 3ω then G has tree-width at least ω , and*
- (ii) *conversely, if G has no $(\omega + 1)$ -connected set of size at least 3ω , then G has tree-width less than 4ω .*

It is worth noting that the theorem proved by Diestel, Jensen, Gorbonov and Thomassen [4] is actually slightly stronger—in (ii), it is not necessary that G have no $(\omega + 1)$ -connected set, but only that G have no *externally* $(\omega + 1)$ -connected set, a stronger property.

Robertson and Seymour [23] proved that high-order tangles give rise to large *grid-minors*; recall that, for $k \in \mathbb{N}$, the $n \times n$ -grid is the graph with vertex set $\{v_{i,j} : i, j \in \{1, \dots, n\}\}$ with $v_{i,j}$ adjacent to $v_{i',j'}$ if and only if $|i - i'| + |j - j'| = 1$. Together with the connection between high-order tangles and highly-connected sets, this implies the following.

Theorem 4.1.6 (Grid Theorem). *For each $n \in \mathbb{N}$, there exists $N \in \mathbb{N}$ such that if G is a graph containing an N -connected set of size at least N , then G contains an $n \times n$ -grid-minor.*

Theorem 4.1.6 was also proved directly by Diestel, Jensen, Gorbonov and Thomassen [4]. Theorem 4.1.6 also follows from the results presented here.

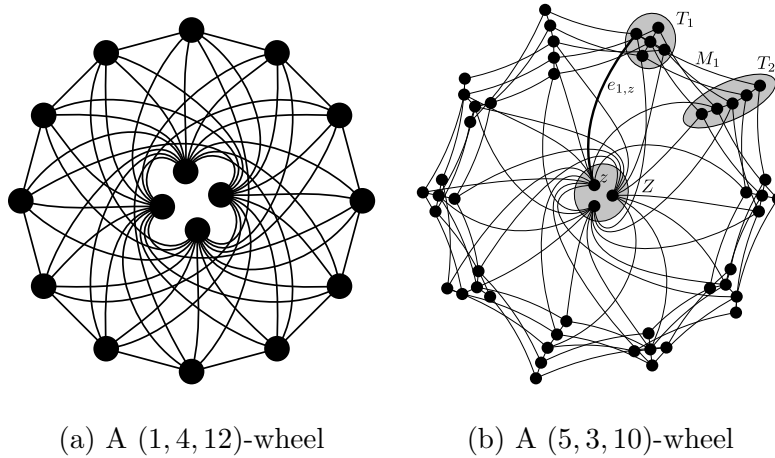


Figure 4.5: Two examples of an (r, ℓ, n) -wheel for different values of r , ℓ and n

4.1.3 The main theorem

Next we describe a graph construction which gives the unavoidable-minors for our main theorem.

For each $r, \ell, n \in \mathbb{N}$ with $r \geq 1$ and $n \geq 3$, an (r, ℓ, n) -wheel is a graph G constructed as follows:

1. let $(T_i : i \in \mathbb{Z}_n)$ be a sequence of vertex-disjoint r -vertex trees indexed by the n -element cyclic group, \mathbb{Z}_n , and let G_1 be the disjoint union of these trees;
2. let G_2 be a graph obtained from G_1 by adding, for each $i \in \mathbb{Z}_n$, an r -edge matching, M_i , between $V(T_i)$ and $V(T_{i+1})$;
3. let G be a graph obtained from G_2 by adding an independent set, Z , of ℓ new vertices such that, for each $z \in Z$ and each $i \in \mathbb{Z}_n$, z is adjacent to exactly one of the vertices in $V(T_i)$, along an edge $e_{i,z}$.

The trees in $(T_i : i \in \mathbb{Z}_n)$ are called the *rim-trees* of the (r, ℓ, n) -wheel. The matchings $(M_i : i \in \mathbb{Z}_n)$ are called the *rim-matchings* of the (r, ℓ, n) -wheel. The vertices in Z are called *hubs* of the (r, ℓ, n) -wheel. The edges in $(e_{i,z} : i \in \mathbb{Z}_n, z \in Z)$ are called *spokes* of the (r, ℓ, n) -wheel. See Figure 4.5. When we do not care about the values of r , ℓ or n , we will call an (r, ℓ, n) -wheel a *generalized wheel*.

We say that an (r, ℓ, n) -wheel has *length* n .

We will prove that generalized wheels, together with complete-bipartite graphs give classes of unavoidable-minors for graphs with large θ -connected sets.

Theorem 4.1.7. *For each $\theta, n \in \mathbb{N}$ with $\theta \geq 2$, there exists $N \in \mathbb{N}$ such that, if G is a graph containing a θ -connected set of size at least N , then either G contains a $K_{\theta, n}$ -minor or there exists $r, \ell \in \mathbb{N}$ with $2r + \ell = \theta$ such that G contains an (r, ℓ, n) -wheel-minor.*

We actually prove Theorem 4.4.1, a stronger version of Theorem 4.1.7.

4.2 Homogeneous wheels

We can further refine the structure of generalized-wheels by applying Ramsey-theoretic techniques; this can be useful in applications.

For $r, \ell, n \in \mathbb{N}$ with $r \geq 1$ and $n \geq 3$, a *homogeneous (r, ℓ, n) -wheel* is a graph constructed as follows:

1. let T be an r -vertex tree, and let G_1 be the disjoint union of n copies of T , named T_1, \dots, T_n , such that, for each $v \in V(T)$ and each $i \in \{1, \dots, n\}$, the copy of v in T_i is labelled v_i ;
2. Let G_2 be obtained from G_1 by adding an edge between v_i and v_{i+1} for each $v \in V(T)$ and each $i \in \{1, \dots, n-1\}$;
3. Let G_3 be obtained from G_2 by adding an arbitrary matching of size r between $\{v_1 : v \in V(T)\}$ and $\{v_n : v \in V(T)\}$;
4. Let G be obtained from G_3 by adding a set, Z of ℓ new vertices, to G_3 , where each $z \in Z$ is adjacent to each copy of some vertex $v_z \in V(T)$.

See Figure 4.6. Note that a homogeneous (r, ℓ, n) -wheel is an (r, ℓ, n) -wheel.

Theorem 4.2.1. *For $\theta, n \in \mathbb{N}$ with $\theta \geq 2$ and $n \geq 3$, there exists $N \in \mathbb{N}$ such that if G is a graph containing a θ -connected set of size at least N , then either G contains a $K_{\theta, n}$ -minor or there exists $r, \ell \in \mathbb{N}$ such that $2r + \ell = \theta$ and G contains a homogeneous (r, ℓ, n) -wheel-minor.*




Figure 4.6: An example of a homogeneous $(4, 3, 12)$ -wheel

We will actually prove Theorem 4.4.2, a stronger version of Theorem 4.2.1.

Applying Theorem 4.2.1 to small values of θ , gives results that are very similar to those discussed in Subsection 4.1.1.

Corollary 4.2.2. *For each $n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that if G is a graph containing a 1-connected set of size at least N , then G contains a minor isomorphic to either a degree- n star or a path of length n .*

Corollary 4.2.3. *For each $n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that if G is a graph containing a 2-connected set of size at least N , then G contains a minor isomorphic to either $K_{2,n}$ or a cycle length n .*

Corollary 4.2.4. *For each $n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that if G is a graph containing a 3-connected set of size at least N , then G contains a minor isomorphic to either $K_{3,n}$ or a wheel of length n .*

Corollary 4.2.5. *For each $n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that if G is a graph containing a 4-connected set of size at least N , then G contains a minor isomorphic to one of the following: $K_{4,n}$, a 2-wheel of length n , a cyclic ladder of length n , or a Möbius ladder of length n .*


Because the vertex set of a θ -connected graph is a θ -connected set, Theorems 4.2.2, 4.2.3, 4.2.4 and 4.2.5 imply Theorems 4.1.1, 4.1.2, 4.1.3 and 4.1.4, respectively.

Theorem 4.2.1 also implies the Grid Theorem (Theorem 4.1.6). This is proved in Subsection 4.10.3. Our proof does not make use of the Grid Theorem, and hence provides an alternative proof of the Grid Theorem.


A *hub* of a cyclic ladder or Möbius ladder is a vertex not in the ladder that is incident with each vertex in one of the two cycles used to construct the ladder; see Figure 4.7 (c) and (d).

Applying Theorem 4.2.1 with $\theta = 5$ shows that each graph with a sufficiently large 5-connected set contains a large homogeneous wheel-minor from one of the classes illustrated in Figure 4.7.


Corollary 4.2.6. *For each $n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that if G is a graph containing a 5-connected set of size at least N , then G contains a minor isomorphic to one of the following: $K_{5,n}$, a 3-wheel of length n , a cyclic ladder of length n with one hub, or a Möbius ladder of length n with one hub.*




(a) $K_{5,n}$



(b) A cycle with 3 hubs



(c) A cyclic-ladder with 1 hub



(d) A Möbius-ladder with 1 hub

Figure 4.7: The unavoidable-minors for graphs with large 5-connected sets.

A *twisted-triple-ladder of length n* is a graph obtained from 3 disjoint cycles of length n with vertices $\{u_1, \dots, u_n\}$, $\{v_1, \dots, v_n\}$ and $\{w_1, \dots, w_n\}$ respectively, by adding, for each $i \in \{1, \dots, n\}$, an edge between u_i and v_i and an edge between v_i and w_i , and replacing the edges $\{u_1u_n, v_1v_n, w_1w_n\}$ with an arbitrary matching between $\{u_1, v_1, w_1\}$ and $\{u_n, v_n, w_n\}$; there are four isomorphism classes of twisted-triple-ladders of length n , depicted in Figure 4.8 (e).

Applying Theorem 4.2.1 with $\theta = 6$ shows that each graph with a sufficiently large 6-connected set contains a large homogeneous wheel minor from one of the classes illustrated in Figure 4.8.

Corollary 4.2.7. *For each $n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that if G is a graph containing a 6-connected set of size at least N , then G contains a minor isomorphic to one of the following: $K_{5,n}$, a 3-wheel of length n , a cyclic ladder of length n with one hub, or a Möbius ladder of length n with one hub.*


4.3 Connectivity and tangles in wheels

It is possible that a generalized wheel graph could be constructed in multiple ways. To avoid ambiguity when referring to the pieces used in the construction of a generalized wheel, we define a *model* of a generalized wheel as follows. For $r, \ell, n \in \mathbb{N}$ with $r \geq 1$ and $n \geq 3$, a *model* for an (r, ℓ, n) -wheel W is a tuple $\mathcal{W} = (\mathbb{T}, \mathbb{M}, Z, \mathfrak{e})$ where

1. $\mathbb{T} = (T_i : i \in \mathbb{Z}_n)$ is a collection of vertex-disjoint, r -vertex trees in W ,
2. $\mathbb{M} = (M_i : i \in \mathbb{Z}_n)$ where, for $i \in \mathbb{Z}_n$, M_i is an r -edge matching between $V(T_i)$ and $V(T_{i+1})$,
3. Z is an ℓ -element subset of $V(W)$ disjoint from T_1, \dots, T_n , and
4. $\mathfrak{e} = (e_{i,z} : i \in \mathbb{Z}_n, z \in Z)$ where, for $i \in \mathbb{Z}_n$ and $z \in Z$, $e_{i,z}$ is an edge of W incident with z and incident with a vertex of T_i .


Given a length n generalized wheel with a model $\mathcal{W} = (\mathbb{T}, \mathbb{M}, Z, \mathfrak{e})$, for $i \in \mathbb{Z}_n$, let T_i^+ be the tree induced by the edge-set $E(T_i) \cup M_i \cup \{e_{i,z} : z \in Z\}$. We call T_i^+ an *augmented tree* of the model \mathcal{W} .

For $r, \ell, n, \theta \in \mathbb{N}$ with $r \geq 1$, $\theta = 2r + \ell$ and $n \geq \theta$, if W is an (r, ℓ, n) -wheel with model \mathcal{W} , then the *fundamental tangle of \mathcal{W}* is the family consisting of the sets $A \subseteq E(W)$ such




(a) $K_{6,n}$

(b) 4-wheel



(c) Cyclic ladder with 2 hubs

(d) Möbius ladder with 2 hubs



(e) Twisted-triple-ladder

Figure 4.8: The unavoidable-minors for graphs with large 6-connected sets.

that $\lambda(A) < \theta$ and A does not contain the edge-set of any augmented tree of \mathcal{W} . We will show in Theorem 4.3.4 that, if $n > \frac{1}{2}(\delta - 1)(\theta - 1)$, then the fundamental tangle of \mathcal{W} is, in fact, a tangle.

Lemma 4.3.1. *For $r, \ell, n, \theta \in \mathbb{N}$ with $r \geq 1$, $\theta = 2r + \ell$ and $n \geq \theta$, if G is an (r, ℓ, n) -wheel with model $\mathcal{W} = (\mathbb{T}, \mathbb{M}, Z, \mathfrak{e})$, and (A, B) is a $(\theta - 1)$ -separation in G , then at most one of A and B contains the edge-set of an augmented tree of \mathcal{W} .*

Proof. If $E(T_i^+) \subseteq A$ and $E(T_{i'}^+) \subseteq B$, then

$$\bigcup_{j \in \mathbb{Z}_n - \{i, i'\}} E(M_j) \cup Z$$

is the union of $2r + \ell$ vertex-disjoint paths that are each incident with an edge in $E(T_i^+) \subseteq A$ and incident with an edge in $E(T_{i'}^+) \subseteq B$, contradicting the fact that (A, B) is a $(2r + \ell - 1)$ -separation. \square

Lemma 4.3.2. *For $r, \ell, n, \theta \in \mathbb{N}$ with $r \geq 1$, $\theta = 2r + \ell$, and $n \geq \theta$, if W is an (r, ℓ, n) -wheel with model $\mathcal{W} = (\mathbb{T}, \mathbb{M}, Z, \mathfrak{e})$, \mathcal{T} is the fundamental tangle of \mathcal{W} , and $A \in \mathcal{T}$, then A contains edges from at most $\theta - 1$ distinct augmented trees of \mathcal{W} .*

Proof. If $i \in \mathbb{Z}_n$ and $A \cap E(T_i^+) \neq \emptyset$, then $E(T_i^+)$ contains an edge in both A and $E(W) - A$; therefore, $E(T_i^+)$ contains a vertex v_i incident with both an edge in $E(T_i^+) \cap A$ and an edge in $E(T_i^+) - A$. The vertex v_i has degree at least two in T_i^+ , so $v_i \in V(T_i)$. For distinct $i, i' \in \mathbb{Z}_n$, with $A \cap E(T_i^+) \neq \emptyset$ and $A \cap E(T_{i'}^+) \neq \emptyset$, $V(T_i) \cap V(T_{i'}) = \emptyset$, so $v_i \neq v_{i'}$. Thus, the number of indices $i \in \mathbb{Z}_n$ such that $A \cap E(T_i^+) \neq \emptyset$ is at most $\lambda(A)$, which is less than θ . \square

Lemma 4.3.3. *For $r, \ell, n, \theta \in \mathbb{N}$ with $r \geq 1$, $\theta = 2r + \ell$ and $n \geq \theta$, if W is an (r, ℓ, n) -wheel with model \mathcal{W} , \mathcal{T} is the fundamental tangle of \mathcal{W} , and $A \in \mathcal{T}$, then $|A| \leq (\theta - 1)(\theta - 2)$.*

Proof. The set A has edges in at most $\theta - 1$ distinct augmented trees of \mathcal{W} by Lemma 4.3.2. Each augmented tree has $\theta - 1$ edges, at least one of which is not in A , giving the desired bound. \square

We now prove that the fundamental tangle of a model is, in fact, a tangle.

Theorem 4.3.4. *For $r, \ell, n, \delta, \theta \in \mathbb{N}$ with $r \geq 1$, $\theta = 2r + \ell$, $\delta \geq 4$ and $n > \frac{1}{2}(\delta - 1)(\theta - 1)$, if W is an (r, ℓ, n) -wheel with model $\mathcal{W} = (\mathbb{T}, \mathbb{M}, Z, \mathfrak{e})$, then the fundamental tangle, \mathcal{T} , of \mathcal{W} is a θ -tangle with covering number at least δ .*

Proof. Tangle properties (T1) and (T5) hold by construction. Tangle property (T2) holds by Lemma 4.3.1.

Let $\{X_1, \dots, X_m\} \subseteq \mathcal{T}$ with $\bigcup_{i=1}^m X_i = E(W)$. By Lemma 4.3.1, for $j \in \{1, \dots, m\}$, X_j contains edges in at most $\theta - 1$ distinct augmented trees of \mathcal{W} by Lemma 4.3.2. For $i \in \mathbb{Z}_n$, the augmented tree T_i^+ contains edges from at least two different sets in $\{X_1, \dots, X_m\}$; also, the augmented trees are edge-disjoint. Therefore, by double counting the number of pairs $(i, j) \in \mathbb{Z}_n \times \{1, \dots, m\}$ such that $X_i \cap E(T_j) \neq \emptyset$,

$$(\delta - 1)(\theta - 1) < 2n \leq (\theta - 1)m,$$

so $m > \delta \geq 4$. Thus, tangle property (T4') holds, and \mathcal{T} is a tangle with covering number at least δ . \square

We also define, for $\theta, n \in \mathbb{N}$ with $n \geq \theta \geq 3$, the *fundamental tangle of $K_{\theta, n}$* to be the family consisting of the sets $A \subseteq E(K_{\theta, n})$ such that $\lambda(A) < \theta$ and there exists $v \in V(K_{\theta, n})$ such that A does not contain any edge incident with v . We will show in Theorem 4.3.6 that, if $n > \frac{3}{2}(\theta - 1)$, then the fundamental tangle of $K_{\theta, n}$ is, in fact, a tangle.

Lemma 4.3.5. *For $\theta, n \in \mathbb{N}$ with $\theta \geq 3$, if \mathcal{T} is the fundamental tangle of $K_{\theta, n}$ and (A, B) is a $(\theta - 1)$ -separation in $K_{\theta, n}$, then either $A \in \mathcal{T}$ or $B \in \mathcal{T}$.*

Proof. Suppose the contrary, so there exist vertices $v_A, v_B \in V(K_{\theta, n})$ such that A contains each edge incident with v_A and B contains each edge incident with v_B . Then v_A and v_B are not adjacent, because an edge between v_A and v_B could not be in A nor in B . Therefore, v_A and v_B have the same set of neighbours, and this common set of neighbours has size at least θ . Each common neighbour of v_A and v_B is incident with both an edge in A and an edge in B , contradicting the fact that (A, B) is a $(\theta - 1)$ -separation. \square

Theorem 4.3.6. *For $n, \theta, \delta \in \mathbb{N}$ with $\theta \geq 3$, $\delta \geq 4$, and $n > \frac{1}{2}(\delta - 1)(\theta - 1)$, if \mathcal{T} is the fundamental tangle of $K_{\theta, n}$, then \mathcal{T} is a tangle with covering-number at least δ .*

Proof. Tangle properties (T1) and (T5) hold by construction. Tangle property (T2) holds by Lemma 4.3.5.

Fix a bipartition (A, B) of $K_{\theta, n}$ such that $|A| = \theta$, $|B| = n$ and each vertex in B has degree θ , and each vertex in A is adjacent to each vertex in B .

Let $\{X_1, \dots, X_m\} \subseteq \mathcal{T}$ with $\bigcup_{i=1}^m X_i = E(K_{\theta, n})$. For $i \in \{1, \dots, m\}$, let $V_i \subseteq B$ be the set of vertices in B incident with an edge in X_i ; note that each vertex in V_i is incident with

both an edge in X_i and an edge in $E(K_{\theta,n}) - X_i$, so

$$|V_i| \leq \lambda(V_i) < \theta.$$

Each vertex $v \in B$ is incident with edges in at least two different sets in $\{X_1, \dots, X_m\}$. Therefore, by double counting the number of pairs $(v, j) \in B \times \{1, \dots, m\}$ such that v is incident with an edge in X_j

$$(\delta - 1)(\theta - 1) < 2n \leq (\theta - 1)m,$$

so $m \geq \delta \geq 4$. Thus, tangle property (T4') holds, and \mathcal{T} is a tangle with covering number at least δ . \square

4.3.1 Weak θ -connectivity

A graph G is *weakly θ -connected* if, for each $t < \theta$ and each t -separation (A, B) of G we have $\min\{|A|, |B|\} < t^2$.

Lemma 4.3.7. *If $r, \ell, n, \theta \in \mathbb{N}$ with $r \geq 1$, $\theta = 2r + \ell$ and $n > 2(\theta - 1)$, then any (r, ℓ, n) -wheel is weakly θ -connected.*

Proof. Let W be an (r, ℓ, n) -wheel with model \mathcal{W} , and let \mathcal{T} be the fundamental tangle of \mathcal{W} . If (A, B) is a $(\theta - 1)$ -separation, then \mathcal{T} contains at least one of A or B by tangle property (T2); without loss of generality, $A \in \mathcal{T}$. By Lemma 4.3.3,

$$|A| \leq (\theta - 1)(\theta - 2) < (\theta - 1)^2.$$

\square

Lemma 4.3.8. *For $\theta, n \in \mathbb{N}$ with $\theta \geq 2$ and $n > 2(\theta - 1)$, $K_{\theta,n}$ is weakly θ -connected.*

Proof. Let \mathcal{T} be the fundamental tangle of $K_{\theta,n}$. If (A, B) is a $(\theta - 1)$ -separation, then \mathcal{T} contains at least one of A or B by tangle property (T2); without loss of generality, $A \in \mathcal{T}$. Each vertex incident with an edge in A is also incident with an edge in $E(K_{\theta,n}) - A$, so there are at most $\lambda(A) < \theta$ such vertices. Therefore,

$$|A| \leq \frac{1}{4}(\theta - 1)^2 < (\theta - 1)^2.$$

\square

Our results can be formulated in terms of unavoidable-minors for weakly θ -connected graphs as follows.

Theorem 4.3.9. *For $\theta, n \in \mathbb{N}$ with $n, \theta \geq 3$, and for $f : \mathbb{N} \rightarrow \mathbb{N}$, there exists $N \in \mathbb{N}$ such that, if G is graph with $|E(G)| \geq N$ and, for each $t < \theta$ and each t -separation (A, B) of G , $\min\{|A|, |B|\} < f(t)$, then G contains a minor G' such that G' is isomorphic to either $K_{\theta, n}$ or an (r, ℓ, n) -wheel, where $2r + \ell = \theta$.*

Theorem 4.3.9 is proved as a consequence of Theorem 4.4.1. It implies the following result for weakly θ -connected graphs.

Corollary 4.3.10. *For $\theta, n \in \mathbb{N}$ with $n, \theta \geq 3$, there exists $N \in \mathbb{N}$ such that if G is a weakly θ -connected graph with $|E(G)| \geq N$, then G contains a minor G' such that G' is isomorphic to either $K_{\theta, n}$ or an (r, ℓ, n) -wheel, where $2r + \ell = \theta$.*

Proof. By Theorem 4.3.9, with $f(t) = t^2$. □

4.3.2 Large θ -connected sets

Large generalized wheels contain large θ -connected sets. For $r, \ell, n \in \mathbb{N}$ with $r \geq 1$ and $n \geq 3$ a *transversal* of a model $\mathcal{W} = (\mathbb{T}, \mathbb{M}, Z, \mathfrak{e})$ for an (r, ℓ, n) -wheel W is a set of vertices, $X \subseteq V(G)$ such that, $X \cap Z = \emptyset$ and, for $i \in \mathbb{Z}_n$, $|X \cap V(T_i)| = 1$.

Lemma 4.3.11. *For $r, \ell, n, \theta \in \mathbb{N}$ with $r \geq 1$, $n \geq 3$ and $\theta = 2r + \ell$, if W is an (r, ℓ, n) -wheel with model \mathcal{W} and X is a transversal for \mathcal{W} , then X is a θ -connected set.*

Proof. Suppose $Y, Z \subseteq X$ with $|Y| = |Z| = t \leq \theta$ and (A, B) is a separation of W such that each vertex in Y is incident with an edge in A and each vertex in Z is incident with an edge in B . Without loss of generality, A is in the fundamental tangle of \mathcal{W} .

Let $i \in \mathbb{Z}_n$ such that $Y \cap V(T_i) \neq \emptyset$. Note that T_i contains a vertex, v_i , incident with an edge in A —the vertex in $Y \cap V(T_i)$. By definition of the fundamental tangle, the augmented tree, there exists $e'_i \in E(T_i^+) \cap B$; each edge in T_i^+ is incident with at least one vertex in $V(T_i)$, so e'_i is incident with some vertex $v'_i \in V(T_i)$. Therefore, the path between v_i and v'_i in T_i contains some vertex, $\hat{v}_i \in V(T_i)$ incident with both an edge in A and an edge in B . Thus, because X is a transversal of \mathcal{W} , the order of the separation (A, B) is at least $|Y|$, so X is a θ -connected set. □

4.4 Alignment of minors and tangles

Recall from Subsection 2.4.1 that, if G is a graph with a minor G' , and \mathcal{T}' is a θ -tangle in G' , then \mathcal{T}' induces the tangle \mathcal{T} in G consisting of the sets $A \subseteq E(G)$ for which $\lambda_G(A) < \theta$ and $A \cap E(G') \in \mathcal{T}'$.

For $r, \ell, n, \delta, \theta \in \mathbb{N}$ with $r \geq 1$, $\theta = 2r + \ell$, $\delta \geq 4$ and $n > \frac{1}{2}(\delta - 1)(\theta - 1)$, if

1. G is a graph,
2. \mathcal{T} is a tangle in G of order at least θ ,
3. G' is a minor of G isomorphic to an (r, ℓ, n) -wheel,
4. \mathcal{W} is a model of G' ,
5. \mathcal{T}' is the fundamental tangle of \mathcal{W} ,
6. $\tilde{\mathcal{T}}$ is the θ -tangle in G induced by \mathcal{T}' , and
7. $\tilde{\mathcal{T}} \subseteq \mathcal{T}$,

then we say that G' is a \mathcal{T} -aligned (r, ℓ, n) -wheel-minor of G . Note that the tangle \mathcal{T} may have higher order than the induced tangle $\tilde{\mathcal{T}}$, but these tangles agree on separations of order less than the order of $\tilde{\mathcal{T}}$.

For $n, \theta, \delta \in \mathbb{N}$ with $\theta \geq 2$, $\delta \geq 4$, and $n > \frac{1}{2}(\delta - 1)(\theta - 1) - \theta$, if

1. G is a graph,
2. \mathcal{T} is a tangle in G of order at least θ ,
3. G' is a minor of G isomorphic to $K_{\theta, n}$,
4. \mathcal{T}' is the fundamental tangle of G' ,
5. $\tilde{\mathcal{T}}$ is the θ -tangle in G induced by \mathcal{T}' , and
6. $\tilde{\mathcal{T}} \subseteq \mathcal{T}$,

then we say that G' is a \mathcal{T} -aligned $K_{\theta, n}$ -minor of G .

We can now state our main theorem in its full generality.

Theorem 4.4.1. *There exists a function $f_{4.4.1} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that for $\theta, n, N \in \mathbb{N}$ with $\theta \geq 2$, $n > 2(\theta - 1)$, and $N = f_{4.4.1}(\theta, n)$, if G is a graph containing a θ -tangle \mathcal{T} with covering-number at least N , then either G contains a \mathcal{T} -aligned $K_{\theta, n}$ -minor or there exists $r, \ell \in \mathbb{N}$ such that $2r + \ell = \theta$ and G contains a \mathcal{T} -aligned (r, ℓ, n) -wheel-minor.*

Theorem 4.4.1 is proved in Section 4.9. Theorem 4.4.1 implies the weaker version of our main theorem, Theorem 4.1.7, restated here.

Theorem 4.1.7. *For each $\theta, n \in \mathbb{N}$ with $\theta \geq 2$, there exists $N \in \mathbb{N}$ such that, if G is a graph containing a θ -connected set of size at least N , then either G contains a $K_{\theta, n}$ -minor or there exists $r, \ell \in \mathbb{N}$ with $2r + \ell = \theta$ such that G contains an (r, ℓ, n) -wheel-minor.*

Proof. Let $N' = f_{4.4.1}(\theta, n)$ and let $N = (\theta - 1)(N' - 1)$. By Theorem 2.9.1, G contains a θ -tangle with covering number at least N . By Theorem 4.4.1, G contains either a $K_{\theta, n}$ -minor or an (r, ℓ, n) -wheel-minor with $2r + \ell = \theta$, as desired. \square

Theorem 4.4.1 also implies the weakly θ -connected version of our main theorem, Theorem 4.3.9, restated here.

Theorem 4.3.9. *For $\theta, n \in \mathbb{N}$ with $n, \theta \geq 3$, and for $f : \mathbb{N} \rightarrow \mathbb{N}$, there exists $N \in \mathbb{N}$ such that, if G is graph with $|E(G)| \geq N$ and, for each $t < \theta$ and each t -separation (A, B) of G , $\min\{|A|, |B|\} < f(t)$, then G contains a minor G' such that G' is isomorphic to either $K_{\theta, n}$ or an (r, ℓ, n) -wheel, where $2r + \ell = \theta$.*

Proof. Let $k = \max\{f(t) : t < \theta\}$, let $N' = \max\{f_{4.4.1}(\theta, n), 4\}$ and let $N = kN'$. Let $\mathcal{T} \subseteq 2^{E(G)}$ be the family consisting of the sets $A \subseteq E(G)$ such that $\lambda_G(A) < \theta$ and $|A| < k$. Note that $|E(G)| > 2k$, so for each $(\theta - 1)$ -separation (A, B) in G , exactly one of A or B has fewer than k edges, so exactly one of A or B is in \mathcal{T} . Therefore, \mathcal{T} satisfies tangle axioms (T1) and (T2). If $A' \subseteq A \in \mathcal{T}$ and $\lambda(A') < \theta$, then $|A'| < |A| < k$, so $A' \in \mathcal{T}$; therefore \mathcal{T} satisfies tangle axiom (T3). Because $N' \geq 4$ and $|E(G)| \geq 4k$, $E(G)$ cannot be partitioned into three sets each of size less than k , so \mathcal{T} satisfies tangle axiom (T4). Also, for $e \in E(G)$, $|E(G) - \{e\}| \geq k$, so \mathcal{T} satisfies tangle axiom (T5). Thus, \mathcal{T} is a tangle of order θ in G . Note that \mathcal{T} has covering number at least N' , so, by Theorem 4.4.1, G contains either a $K_{\theta, n}$ -minor or an (r, ℓ, n) -wheel-minor with $2r + \ell = \theta$, as desired. \square

We will also prove a similar theorem for homogeneous wheels:

Theorem 4.4.2. *There exists a function $f_{4.4.2} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that, for each $\theta, n, N \in \mathbb{N}$, with $\theta > 1$, $n > 2(\theta - 1)$ and $N = f_{4.4.2}(\theta, n)$, if G is a graph containing a θ -tangle \mathcal{T} with covering-number at least N , then either G contains a \mathcal{T} -aligned $K_{\theta, n}$ -minor or there exists $r, \ell \in \mathbb{N}$ such that $2r + \ell = \theta$ and G contains a \mathcal{T} -aligned, homogeneous (r, ℓ, n) -wheel-minor.*

Theorem 4.4.2 is proved in Section 4.10. Theorem 4.4.2 implies the weaker version of our main theorem, Theorem 4.2.1, restated here.

Theorem 4.2.1. *For $\theta, n \in \mathbb{N}$ with $\theta \geq 2$ and $n \geq 3$, there exists $N \in \mathbb{N}$ such that if G is a graph containing a θ -connected set of size at least N , then either G contains a $K_{\theta, n}$ -minor or there exists $r, \ell \in \mathbb{N}$ such that $2r + \ell = \theta$ and G contains a homogeneous (r, ℓ, n) -wheel-minor.*

Proof. Let $N' = f_{4.4.2}(\theta, n)$ and let $N = (\theta - 1)(N' - 1)$. By Theorem 2.9.1, G contains a θ -tangle with covering number at least N . By Theorem 4.4.1, G contains either a $K_{\theta, n}$ -minor or an homogeneous (r, ℓ, n) -wheel-minor with $2r + \ell = \theta$, as desired. \square

The grid theorem also has a stronger formulation in terms of \mathcal{T} -aligned minors. If G is the $n \times n$ -grid with $V(G) = \{v_{i,j} : i, j \in \{1, \dots, n\}\}$ and $v_{i,j}$ adjacent to $v_{i',j'}$ if and only if $|i - i'| + |j - j'| = 1$, then, for $i \in \{1, \dots, n\}$, row i of G is the subgraph of G induced by the vertex set $\{v_{i,j} : j \in \{1, \dots, n\}\}$. The *fundamental tangle of the $n \times n$ -grid* is the family consisting of the sets $A \subseteq E(G)$ for which $\lambda_G(A) < n$ and A does not contain the edge set of any row of G . Robertson and Seymour [28] proved that the fundamental tangle of the $n \times n$ -grid is a tangle of order n , and proved that each high-order tangle, \mathcal{T} in a graph gives rise to a large, \mathcal{T} -aligned grid-minor.

Theorem 4.4.3. (Grid Theorem). *For each $n \in \mathbb{N}$, there exists some $N \in \mathbb{N}$ such that if G contains a tangle \mathcal{T} of order at least N , then G contains a \mathcal{T} -aligned $n \times n$ -grid-minor.*

We will show in Section 4.10 that Theorem 4.4.3 follows from Theorem 4.4.1.

4.5 Necklace

Working directly with wheel-minors is unwieldy, so this section introduces a closely related structure, called a *necklace*, to prove Theorem 4.4.1.




Figure 4.9: (B, X, Y) is an example of a 4-bead— B is connected and for each $X' \subseteq X$ and each $Y' \subseteq Y$ with $|X'| = |Y'| \leq 4$, B contains $|X'|$ vertex-disjoint paths between X' and Y' ; a collection of 4 vertex-disjoint paths between a particular example of X' and Y' is shown by the bold edges.

If X and Y are two sets of vertices in G , then an (X, Y) -path in G is a path P such that either P contains only one vertex, which is in $X \cap Y$, or P has length greater than 0 and one of the degree-1 vertices in P is in X and the other is in Y ; we will also call an (X, Y) -path a *path between X and Y* .

For $t \in \mathbb{N}$ with $t \geq 1$, a t -bead is a triple (B, X, Y) where B is a connected graph and $X, Y \subseteq V(B)$ such that $\min\{|X|, |Y|\} \geq t$ and for any $X' \subseteq X$ and any $Y' \subseteq Y$ with $|X'| = |Y'| \leq t$, B contains $|X'|$ vertex-disjoint (X', Y') -paths. A 0 -bead is a triple (B, X, Y) where B is a graph and $X, Y \subseteq V(B)$ such that B is the union of two connected subgraphs, B^1 and B^2 with $X \subseteq V(B^1)$, $Y \subseteq V(B^2)$. For example, if B is a complete bipartite graph with bipartition (X, Y) and $\min\{|X|, |Y|\} \geq t$, then (B, X, Y) is a t -bead. A grid is another example of a bead—if B is an $n \times n$ -grid, X is the set of vertices in the leftmost column, and Y is the set of vertices in the rightmost column, and $n \geq t$, then (B, X, Y) is a t -bead. Figure 4.9 shows an example of a 5-bead.

For $t, s, n \in \mathbb{N}$ with $t \geq \max\{1, s\}$, and $n \geq 3$, if G is a graph, then a (t, s, n) -necklace in G is a cyclic sequence $(B_i : i \in \mathbb{Z}_n)$ of subgraphs of G satisfying the following properties:

- (N1) for each $i \in \mathbb{Z}_n$, $|V(B_i) \cap V(B_{i+1})| = t$;
- (N2) for each $i, j \in \mathbb{Z}_n$ if $j - i \notin \{-1, 0, 1\}$ then $V(B_i) \cap V(B_j) = \emptyset$;
- (N3) for each $i \in \mathbb{Z}_n - \{0\}$, $(B_i, V(B_{i-1}) \cap V(B_i), V(B_i) \cap V(B_{i+1}))$ is a t -bead; and

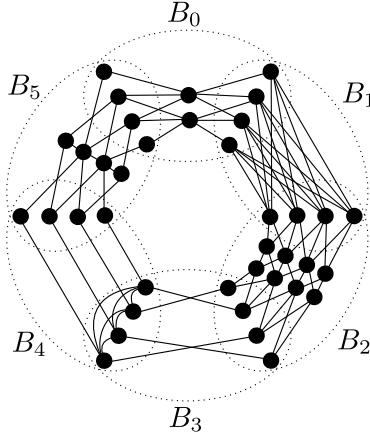


Figure 4.10: $(B_i : i \in \mathbb{Z}_5)$ is an example of a $(4, 2, 6)$ -necklace. $(B_0, V(B_5) \cap V(B_0), V(B_0) \cap V(B_1))$ is a 2-bead and for $i \in \{1, 2, 3, 4\}$, $(B_i, V(B_i), V(B_{i+1}))$ is a 4-bead. Observe that the beads in a necklace need not be edge-disjoint— B_3 and B_4 share edges.

(N4) $(B_0, V(B_{n-1}) \cap V(B_0), V(B_0) \cap V(B_1))$ is an s -bead.

See Figure 4.10. We call B_0 the *weak bead* of the necklace.

A *hub* of a necklace is a vertex $z \in V(G) - \bigcup_{i=1}^n V(B_i)$ such that for each $i \in \mathbb{Z}_n$, z is adjacent to a vertex in $V(B_i)$.

We define the *fundamental tangle* of a necklace in a way similar to the fundamental tangle of a model for a wheel. For each $t, s \in \mathbb{N}$ with $t \geq s$ and $t \geq 1$, if $(B_i : i \in \mathbb{Z}_n)$ is a (t, s, n) -necklace, then the *fundamental tangle* of $(B_i : i \in \mathbb{Z}_n)$ is the family \mathcal{T} consisting of the sets $A \subseteq E(\bigcup_{i \in \mathbb{Z}_n} E(B_i))$ such that $\lambda(A) < t + s$ and, for each $i \in \mathbb{Z}_n$, $E(B_i) \not\subseteq A$.

Lemma 4.5.1. *For each $t, s \in \mathbb{N}$ with $t \geq s$, $t \geq 1$ and $n > 3(t + s)$, if $(B_i : i \in \mathbb{Z}_n)$ is a (t, s, n) -necklace, and \mathcal{T} is the fundamental tangle of $(B_i : i \in \mathbb{Z}_n)$, then \mathcal{T} is a tangle of order $t + s$.*

Proof. Let $G = \bigcup_{i \in \mathbb{Z}_n} E(B_i)$. As was the case for fundamental tangles of wheel models, if (X, Y) is a $(t + s - 1)$ separation in G , then there exists a bead B_i such that either $E(B_i) \subseteq X$ or $E(B_i) \subseteq Y$, but not both; without loss of generality, the latter occurs, so $X \in \mathcal{T}$. Then X contains edges in at most $t + s$ distinct beads, so no three sets in \mathcal{T} can cover S . Therefore, \mathcal{T} is a tangle of order $t + s$. \square

If G is a graph with a tangle \mathcal{T} of order at least $(t + s)$, we say that a (t, s, n) -necklace $(B_i : i \in \mathbb{Z}_n)$ in G is \mathcal{T} -aligned if the fundamental tangle of the necklace induces a subset of \mathcal{T} in G .

We say that a (t, s, n) -necklace is a $(t + s)$ -linked necklace.

To prove Theorem 4.4.1, we first prove Lemma 4.5.2 which shows that a graph with a sufficiently large θ -tangle contains a large $(\theta - \ell)$ -linked necklace with ℓ hubs. Then, in Section 4.9, we will show how to turn the necklace with hubs into a large θ -connected-wheel-minor.

Lemma 4.5.2. *There exists a function $f_{4.5.2} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that for each $\theta, m, n \in \mathbb{N}$ with $\theta \geq 3$, $n \geq 3$ and $m = f_{4.5.2}(\theta, n)$, if G is a graph containing a θ -tangle \mathcal{T} with covering number at least m , then either G contains a \mathcal{T} -aligned $K_{\theta, n}$ -minor or there exists $t, s, \ell \in \mathbb{N}$ with $t \geq s$, $t \geq 1$ and $t + s + \ell = \theta$ such that G contains a \mathcal{T} -aligned (t, s, n) -necklace with ℓ hubs.*

The proof of Lemma 4.5.2 essentially goes by induction on the connectivity θ . In Section 4.6 we find either a $K_{\theta, n}$ -minor (in which case Lemma 4.5.2 holds) or a $(1, 0, n)$ -necklace aligned with our tangle along with θ vertex-disjoint paths between each pair of beads in the $(1, 0, n)$ -necklace. Then, in Section 4.7 and Section 4.8 we show how to use these paths between the beads to increase the necklace connectivity.

If $(B_j : j \in \mathbb{Z}_m)$ and $(B'_i : i \in \mathbb{Z}_n)$ are two necklaces such that, for each $i \in \{1, \dots, n\}$ there exists $j \in \{1, \dots, m\}$ such that $B_j \subseteq B'_i$, then we say that $(B'_i : i \in \mathbb{Z}_n)$ is *supported* by $(B_j : j \in \mathbb{Z}_m)$. Supported necklaces have some crucial properties: if $(B_j : j \in \mathbb{Z}_m)$ is \mathcal{T} -aligned, for some tangle \mathcal{T} in G , then any necklace supported by $(B_j : j \in \mathbb{Z}_m)$ is also \mathcal{T} -aligned. Also, if G contains θ vertex-disjoint paths between any two beads of $(B_j : j \in \mathbb{Z}_m)$, then the same is true of any necklace supported by $(B_j : j \in \mathbb{Z}_m)$. Finally, if G contains a hub z of $(B_j : j \in \mathbb{Z}_m)$, then z is also a hub of any necklace that does not contain z and is supported by $(B_j : j \in \mathbb{Z}_m)$.

4.5.1 Basic properties of beads and necklaces

The following simple observations about beads and necklaces are used throughout the rest of the chapter. First, observe that a $(t + 1)$ -bead is always also a t -bead.

Lemma 4.5.3. *For $t \in \mathbb{N}$, if (B, X, Y) is a $(t + 1)$ -bead, then (B, X, Y) is a t -bead.*

Proof. Suppose first that $t = 0$, so (B, X, Y) is a 1-bead. Then B is connected, so it is the union of two connected subgraphs that contain X and Y : two copies of itself.

Suppose now that $t > 0$. Then B is connected and $\min\{|X|, |Y|\} \geq t + 1 \geq t$. For each $X' \subseteq X$ and $Y' \subseteq Y$ with $|X'| = |Y'| \leq t \leq t + 1$, B contains $|X'|$ vertex-disjoint paths between X' and Y' . \square

Lemma 4.5.4 and Lemma 4.5.5 show that the union of beads is a bead, provided they overlap in the appropriate way.

Lemma 4.5.4. *For $t_1, t_2 \in \mathbb{N}$ with $\max\{t_1, t_2\} \geq 1$, if (B_1, X, Y) is a t_1 -bead, (B_2, Y, Z) is a t_2 -bead, and $V(B_1) \cap V(B_2) = Y$, then $(B_1 \cup B_2, X, Z)$ is a $\min\{t_1, t_2\}$ -bead.*

Proof. By symmetry, we may assume $t_1 \leq t_2$ and $t_2 \geq 1$.

Suppose first that $t_1 = 0$. Then B_1 is the union of two connected subgraphs, B_1^1 and B_1^2 such that $X \subseteq V(B_1^1)$ and $Y \subseteq V(B_1^2)$. Because (B_2, Y, Z) is a t_2 -bead and $t_2 \geq 1$, $Y \neq \emptyset$, so $B_1^2 \cup B_2$ is connected. Therefore, $B_1 \cup B_2$ is the union of two connected subgraphs, B_1^1 and $B_1^2 \cup B_2$ and $X \subseteq V(B_1^1)$ and $Z \subseteq V(B_1^2 \cup B_2)$, so $(B_1 \cup B_2, X, Z)$ is a 0-bead, as desired.

Now suppose $t_1 > 0$, so both B_1 and B_2 are connected and $V(B_1) \cap V(B_2) = Y \neq \emptyset$. Therefore, $B_1 \cup B_2$ is connected. Let $X' \subseteq X$ and $Z' \subseteq Z$ such that $|X'| = |Z'| \leq t_1$. Choose a collection $\{P_x : x \in X'\}$ of vertex-disjoint (X', Y) -paths in B_1 minimizing $|\bigcup_{x \in X'} E(P_x)|$. By minimality, for each $x \in X'$, P_x contains exactly one vertex, y_x , in Y , and y_x is an endpoint of P_x . Because $V(B_1) \cap V(B_2) = Y$, $V(P_x) \cap V(B_2) = V(P_x) \cap Y = \{y_x\}$. Let $Y' = \{y_x : x \in X'\}$. Choose collection $\{Q_y : y \in Y'\}$ of vertex-disjoint (Y', Z') -paths in B_2 . Then $\{P_x \cup Q_{y_x} : x \in X'\}$ is a collection of vertex-disjoint (X', Z') -paths in $B_1 \cup B_2$, as desired. \square

Lemma 4.5.5. *For each $n \in \mathbb{N}$ and each sequence (t_0, \dots, t_{n+1}) of natural numbers with $|\{i \in \{1, \dots, n\} : t_i = 0\}| \leq 1$, if (B_0, \dots, B_{n+1}) is sequence of subgraphs in a graph G such that,*

(i) *for each $i \in \{0, \dots, n\}$, $|V(B_i) \cap V(B_{i+1})| \geq \max\{t_i, t_{i+1}\}$,*

(ii) *for each $i, j \in \{0, \dots, n+1\}$, if $j - i > 1$, then $V(B_i) \cap V(B_j) = \emptyset$,*

(iii) *for each $i \in \{1, \dots, n\}$, $(B_i, V(B_{i-1}) \cap V(B_i), V(B_i) \cap V(B_{i+1}))$ is a t_i -bead*

then

$$\left(\bigcup_{i=1}^n B_i, V(B_0) \cap V(B_1), V(B_n) \cap V(B_{n+1}) \right)$$

is a $(\min\{t_1, \dots, t_n\})$ -bead.

Proof. The proof goes by induction on n . If $n \leq 2$, the lemma follows immediately from Lemma 4.5.4, so suppose $n > 2$. By the induction hypothesis,

$$\left(\bigcup_{i=1}^{n-1} B_i, V(B_0) \cap V(B_1), V(B_{n-1}) \cap V(B_n) \right)$$

is a $(\min\{t_1, \dots, t_{n-1}\})$ -bead. Because $|\{i \in \{1, \dots, n\} : t_i = 0\}| \leq 1$, it cannot be the case that both $\min\{t_1, \dots, t_{n-1}\} = 0$ and $t_n = 0$. Therefore, by Lemma 4.5.4,

$$\left(\bigcup_{i=1}^n B_i, V(B_0) \cap V(B_1), V(B_n) \cap V(B_{n+1}) \right)$$

is a $(\min\{t_1, \dots, t_n\})$ -bead. □

Beads in a necklace always overlap as required in Lemma 4.5.5, so it is easy to prove that taking the union of adjacent beads in a necklace, yields a necklace.

Lemma 4.5.6. *For each $t, s, m, n \in \mathbb{N}$ with $t \geq \max\{s, 1\}$ and $m \geq n \geq 3$, if G is a graph containing a (t, s, m) -necklace $(B_i : i \in \mathbb{Z}_m)$ and $(a_i : i \in \{1, \dots, m\})$ is a sequence of integers such that $0 \leq a_0 < \dots < a_{n-1} \leq m$, then $(B'_i : i \in \mathbb{Z}_n)$ is a (t, s, n) -necklace supported by $(B_i : i \in \mathbb{Z}_m)$ where for each $i \in \mathbb{Z}_n$ $B'_i = \bigcup_{j=a_{i-1}+1}^{a_i} B_j$.*

Proof. For each $i \in \mathbb{Z}_n$,

$$|V(B'_i) \cap V(B'_{i+1})| = |V(B_{a_i}) \cap V(B_{a_{i+1}})| = t,$$

so necklace axiom (N1) holds.

For each $i, j \in \mathbb{Z}_n$ with $j - i \notin \{-1, 0, 1\}$, and each $i', j' \in \mathbb{Z}_m$ such that $B_{i'} \subseteq B'_i$ and $B_{j'} \subseteq B'_j$, $j' - i' \notin \{-1, 0, 1\}$, so $V(B_{i'}) \cap V(B_{j'}) = \emptyset$, so axiom (N2) holds.

By Lemma 4.5.5, for each $i \in \{1, \dots, n-1\}$, $(B'_i, V(B'_{i-1}) \cap V(B'_i), V(B'_i) \cap V(B'_{i+1}))$ is a t -bead, so axiom (N3) holds.

By Lemma 4.5.5, $(B'_0, V(B'_{n-1}) \cap V(B'_0), V(B'_0) \cap V(B'_1))$ is an s -bead, so axiom (N4) holds. □

As a consequence, the length of a necklace can always be reduced.

Corollary 4.5.7. *For $t, s, m, n \in \mathbb{N}$ with $t \geq \max\{s, 1\}$ and $m \geq n \geq 3$, if G is a graph containing a (t, s, m) -necklace (B_1, \dots, B_m) , then G contains a (t, s, n) -necklace supported by (B_1, \dots, B_m) .*

4.6 Initial necklace

In this section we will prove the basis for the inductive proof of Lemma 4.5.2. In particular, we will show that a large θ -tangle either gives rise to a $(1, 0, n)$ -necklace aligned with that tangle along with θ vertex-disjoint paths between each bead, or gives rise a $K_{\theta, n}$ -minor aligned with the tangle.

This is done in two parts: in Subsection 4.6.1, we use the tangle to obtain a large collection of disjoint, connected subgraphs (or a $K_{\theta, n}$ -minor), and in Subsection 4.6.2, these disjoint, connected subgraphs are turned into a $(1, 0, n)$ -necklace (or a $K_{\theta, n}$ -minor).

We start with a few definitions used in this and later sections.

As with wheels and necklaces, a *hub* of a family \mathcal{H} of subgraphs is a vertex that is not in any graph in \mathcal{H} , but has a neighbour in each graph in \mathcal{H} .

If G is a graph and $X \subseteq V(G)$, then a *bridge* B of X in G (or simply an X -*bridge*) is a maximal subgraph of G such that, for any two edges $e, f \in E(B)$, either

1. there exists a cycle $C \subseteq B$ such that $e, f \in E(C)$ and $|V(C) \cap X| \leq 1$, or
2. there exists a path $P \subseteq B$ with endpoints x and y such that $e, f \in E(P)$ and $V(P) \cap X \subseteq \{x, y\}$.

We call $V(B) \cap X$ the *attachment vertices* of the bridge B . Note that each bridge B of X is either a single edge with both ends in X or is a connected component of $G - X$ together with all of the edges between that connected component and X .

4.6.1 From a tangle to disjoint connected subgraphs

In this subsection, we will show that a large tangle gives rise to either a large collection of disjoint, connected subgraphs, none of which are contained in any set in the tangle, or a large $K_{\theta, n}$ -minor aligned with the tangle; see Lemma 4.6.2.

The following lemma proves the intuitive fact that, for any θ -tangle, the deletion of fewer than θ vertices leaves a single connected component that is “large” relative to the tangle, and all of the other connected components lie in a single small set of the tangle, separated from the “large” part by the deleted vertices.

Lemma 4.6.1. *If G is a graph, \mathcal{T} is a θ -tangle in G , and $X \subseteq V(G)$ such that $|X| < \theta$, then G contains a unique X -bridge B_0 such that $E(G) - E(B_0) \in \mathcal{T}$.*

Proof. If $A_1, A_2 \in \mathcal{T}$ and $V(A_i) \cap V(E(G) - A_i) \subseteq X$ for $i \in \{1, 2\}$, then

$$V(A_1 \cup A_2) \cap V(E(G) - (A_1 \cup A_2)) \subseteq X.$$

Therefore, $\lambda(A_1 \cup A_2) \leq |X| < \theta$, so $A_1 \cup A_2 \in \mathcal{T}$.

Let \mathcal{B} be the collection all X -bridges B in G for which $E(B) \in \mathcal{T}$. Note that, for each $B \in \mathcal{B}$,

$$V(B) \cap V(E(G) - E(B)) \subseteq X.$$

It follows by induction on $|\mathcal{B}'|$, for each $\mathcal{B}' \subseteq \mathcal{B}$, $E(\bigcup \mathcal{B}') \in \mathcal{T}$. But $E(G) \notin \mathcal{T}$, so there is some X -bridge B_0 such that $E(B_0) \notin \mathcal{T}$. Because $\lambda_G(E(B_0)) \leq |X| < \theta$, this implies $E(G) - E(B_0) \in \mathcal{T}$.

B_0 is the only such bridge because, for any X -bridge B with $B \neq B_0$, $E(B) \subseteq E(G) - E(B_0)$, so $E(B) \in \mathcal{T}$. \square

Now we can prove that a large tangle either gives rise to a large collection of disjoint, connected subgraphs, none of which are contained in any set in the tangle, or gives rise to a large $K_{\theta, n}$ -minor aligned with the tangle.

For each $m \in \mathbb{N}$, a set A of edges of a graph G is m -covered by a tangle \mathcal{T} in G if there exists $X_1, \dots, X_m \in \mathcal{T}$ such that $A \subseteq \bigcup_{i=1}^m X_m$. A subgraph H of G is m -covered by a tangle \mathcal{T} in G if $E(H)$ is m -covered by \mathcal{T} .

Lemma 4.6.2. *There exists a function $f_{4.6.2} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that, for $\theta, m, n \in \mathbb{N}$ with $\theta > 2$ and $m = f_{4.6.2}(\theta, n)$, if G is a graph and \mathcal{T} is a θ -tangle in G and covering number at least m , then either G contains a collection \mathcal{H} of disjoint, connected subgraphs, none of which are 1-covered by \mathcal{T} , with $|\mathcal{H}| \geq n$, or G contains a \mathcal{T} -aligned $K_{\theta, n}$ -minor.*

Proof. Let

$$\begin{aligned} m' &= 2(\theta - 1)^2(n - 1) + 1, \\ m_1 &= m' \binom{2(\theta - 1)(n - 1)}{\theta} \max\{n - 1, 2(\theta - 1)\}, \\ m_2 &= \binom{2(\theta - 1)(n - 1)}{\theta - 1}, \\ m_3 &= \binom{n(\theta - 1)}{2}, \\ f_{4.6.2}(\theta, n) = m &= m_1 + m_2 + m_3 + n + 1, \end{aligned}$$

and let G be a graph with a θ -tangle \mathcal{T} and covering number at least m .

Let \mathcal{H} be a collection of disjoint, connected subgraphs of G such that, for each $H \in \mathcal{H}$, H is 2-covered but not 1-covered by \mathcal{T} , and $|\mathcal{H}|$ is maximum. If $|\mathcal{H}| \geq n$, then the lemma holds, so we may assume $|\mathcal{H}| < n$. For each $H \in \mathcal{H}$, let $A_1^H, A_2^H \in \mathcal{T}$ such that $E(H) \subseteq A_1^H \cup A_2^H$. Let $G' = G - V(\bigcup_{H \in \mathcal{H}} (A_1^H \cup A_2^H))$.

Claim 4.6.2.1. *Each connected subgraph of G' is 1-covered by \mathcal{T} .*

Proof of Claim. Suppose the contrary, and let H be a minimal connected subgraph of G' that is not 1-covered by \mathcal{T} . Let T be a spanning tree of H , and let $e \in E(T)$ be an edge incident with a leaf v . Let H' be obtained from H by deleting e and, if v is not incident with any other edge in H , deleting v . Then H' is connected and is a proper subgraph of H , so, by minimality of H , H' is 1-covered by \mathcal{T} . Let $A \in \mathcal{T}$ such that $E(H') \subseteq A$. Then $\{A, \{e\}\}$ is a cover of H by two sets in \mathcal{T} , so H is 2-covered by \mathcal{T} , so $\mathcal{H} \cup \{H\}$ satisfies all the requirements of \mathcal{H} , contradicting maximality of $|\mathcal{H}|$. \square (Claim)

Given any set $E' \subseteq E(G)$, let $\partial(E') = V(E') \cap V(E(G) - E')$. Similarly, for any subgraph H of G , let $\partial(H) = \partial(E(H))$. Let $V_\partial = \bigcup_{H \in \mathcal{H}} \partial(A_1^H \cup A_2^H)$, and note that $|V_\partial| \leq 2(\theta - 1)|\mathcal{H}| \leq 2(\theta - 1)(n - 1)$.

Claim 4.6.2.2. *For any V_∂ -bridge B in G , B is m' -covered by \mathcal{T} .*

Proof of Claim. By Claim 4.6.2.1, the connected component C of G' contained in B is 1-covered by \mathcal{T} , so there exists $A \in \mathcal{T}$ such that $E(C) \subseteq A$. Let $X = V(C) \cap \partial(A)$, and note that, because $A \in \mathcal{T}$, $|X| \leq |\partial(A)| \leq \theta - 1$. If $e \in E(B) - A$, then $e \notin E(C)$, so e is incident with a vertex in V_∂ and a vertex in X ; hence, there are at most $|V_\partial||X| \leq 2(\theta - 1)^2(n - 1)$ parallel classes of such edges, and each parallel class is in \mathcal{T} . Therefore, B can be covered by at most $2(\theta - 1)^2(n - 1) + 1 = m'$ sets in \mathcal{T} , so B is m' -covered by \mathcal{T} , proving the claim. \square (Claim)

Let \mathcal{B}_1 denote the set of all V_∂ -bridges B such that $|\partial(B)| \geq \theta$; let \mathcal{B}_2 denote the set of all V_∂ -bridges B such that $|\partial(B)| < \theta$.

Claim 4.6.2.3. *If $\bigcup \mathcal{B}_1$ is not m_1 -covered by \mathcal{T} , then G contains a \mathcal{T} -aligned $K_{\theta, n}$ -minor.*

Proof of Claim. Suppose $\bigcup \mathcal{B}_1$ is not m_1 -covered by \mathcal{T} . For $Z \subseteq V_\partial$ with $|Z| = \theta$, let $\mathcal{B}_1^Z = \{B \in \mathcal{B}_1 : Z \subseteq \partial(B)\}$. For each $B \in \mathcal{B}_1$, there is some $Z \subseteq \partial(B)$ with $|Z| = \theta$, so $\bigcup_{Z \subseteq V_\partial: |Z| = \theta} \mathcal{B}_1^Z = \mathcal{B}_1$. Therefore, by the pigeon-hole principle, there is some $Z \subseteq V_\partial$ with $|Z| = \theta$ such that $\bigcup \mathcal{B}_1^Z$ is not $\left(m_1 \binom{|V_\partial|}{\theta}^{-1}\right)$ -covered by \mathcal{T} .

Because each $B \in \mathcal{B}_1$ is m' -covered by \mathcal{T} , Claim 4.6.2.2 implies

$$|\mathcal{B}_1^Z| > m_1 \left(\binom{|V_\partial|}{\theta} m' \right)^{-1} \geq m_1 \left(\binom{2(\theta-1)(n-1)}{\theta} m' \right) \geq n-1,$$

so $|\mathcal{B}_1^Z| \geq n$. For each $z \in Z$ and $B \in \mathcal{B}_1^Z$, there is some edge $e_{z,B}$ in G between z and B . Let G' be obtained from G by contracting $B - V_\partial$ to a single vertex, v_B for each $B \in \mathcal{B}_1^Z$, and deleting every other edge except $\{e_{z,B} : z \in Z, B \in \mathcal{B}_1^Z\}$.

Note that $Z \subseteq V(G')$ because $Z \subseteq V(B)$ for each $B \in \mathcal{B}_1^Z$. Note also that, for each $B \in \mathcal{B}_1^Z$ and each $z \in Z$, $e_{z,B}$ is an edge from z to v_B in G' , so G' is isomorphic to $K_{\theta, |\mathcal{B}_1^Z|}$. Therefore, we need only show that G' is a \mathcal{T} -aligned minor.

Let (X, Y) be a $(\theta - 1)$ -separation in G . Not every vertex in Z can be incident with both an edge in X and an edge in Y because $|Z| = \theta$, so, without loss of generality, there is some $z_0 \in Z$ such that each edge incident with z_0 in G is in Y . Therefore, each bridge $B \in \mathcal{B}_1^Z$ contains an edge in Y , so at most $\theta - 1$ of these bridges can contain an edge in X .

Note that the bridges in \mathcal{B}_1^Z containing an edge in X can be covered by at most $m'(\theta - 1) < m_1 \binom{|V_\partial|}{\theta}^{-1}$ sets in \mathcal{T} . If $Y \in \mathcal{T}$, then $\bigcup \mathcal{B}_1^Z$ can be covered by at most $m_1 \binom{|V_\partial|}{\theta}^{-1}$ sets in \mathcal{T} , contradicting the choice of Z , so it must be that $X \in \mathcal{T}$.

For each $B \in \mathcal{B}_1^Z$, $e_{z_0, B} \in Y$, so $|\{e_{z,B} : z \in Z\} \cap X| \leq \theta - 1$, and, as we just showed, $\{e_{z,B} : z \in Z\} \cap X = \emptyset$ for all but at most $\theta - 1$ bridges $B \in \mathcal{B}_1^Z$, so

$$|X \cap E(G')| = |X \cap \{e_{z,B} : z \in Z, B \in \mathcal{B}_1^Z\}| \leq (\theta - 1)^2.$$

Note that

$$|\mathcal{B}_1^Z| > m_1 \left(\binom{|V_\partial|}{\theta} m' \right)^{-1} \geq m_1 \left(\binom{2(\theta-1)(n-1)}{\theta} m' \right) \geq 2(\theta - 1),$$

so

$$|E(G')| = |\mathcal{B}_1^Z| |Z| \geq 2(\theta - 1)\theta > 2(\theta - 1)^2 \geq 2|X \cap E(G')|.$$

Thus, for any $(\theta - 1)$ -separation in G , the side in \mathcal{T} contains fewer than half of the edges of G' , so G' is a \mathcal{T} -aligned minor, proving the claim. \square (Claim)

Claim 4.6.2.4. $\bigcup \mathcal{B}_2$ is m_2 -covered by \mathcal{T} .

Proof of Claim. For $Z \subseteq V_\partial$ with $|Z| = \min\{\theta - 1, |V_\partial|\}$, let $\mathcal{B}_2^Z = \{B \in \mathcal{B}_2 : \partial(B) \subseteq Z\}$. For each $B \in \mathcal{B}_2^Z$, $\lambda_G(E(B)) \leq |Z| < \theta$, so either $E(B) \in \mathcal{T}$ or $E(G) - E(B) \in \mathcal{T}$. By

Claim 4.6.2.2, $E(B)$ is m' -covered by \mathcal{T} , and \mathcal{T} has covering number at least $m > m' + 1$, so $E(G) - E(B) \notin \mathcal{T}$, so $E(B) \in \mathcal{T}$. Thus, $E(B) \in \mathcal{T}$ for each $B \in \mathcal{B}_2^Z$. It follows by a simple inductive argument that $\bigcup_{B \in \mathcal{B}_2^Z} E(B) \in \mathcal{T}$.

But, for each $B \in \mathcal{B}_2$, there is some $Z \subseteq V_\partial$ with $\partial(B) \subseteq Z$ and $|Z| = \min\{\theta - 1, |V_\partial|\}$, so $E(B) \subseteq \bigcup_{B' \in \mathcal{B}_2^Z} E(B')$. Therefore,

$$\bigcup \mathcal{B}_2 = \bigcup_{\substack{Z \subseteq V_\partial \\ |Z| = \min\{\theta - 1, |V_\partial|\}}} \bigcup \mathcal{B}_2^Z,$$

so $\bigcup \mathcal{B}_2$ can be covered by at most

$$\binom{|V_\partial|}{\min\{\theta - 1, |V_\partial|\}} \leq \binom{2(\theta - 1)(n - 1)}{\theta - 1} = m_2$$

sets in \mathcal{T} , proving the claim. □ (Claim)

Let E_∂ denote the set of edges of G with both ends in V_∂ .

Claim 4.6.2.5. E_∂ is m_3 -covered by \mathcal{T} .

Proof of Claim. E_∂ consists of at most $\binom{|V_\partial|}{2} \leq \binom{n(\theta - 1)}{2} = m_3$ parallel classes, each of which is in \mathcal{T} . □ (Claim)

If $|\mathcal{A}| \leq n$ and G contains no \mathcal{T} -aligned $K_{\theta, n}$ -minor, then each edge in G is either in $\bigcup \mathcal{A}$, which is n -covered by \mathcal{T} , or $\bigcup_{B \in \mathcal{B}_1} E(B)$, which is m_1 -covered by \mathcal{T} , or $\bigcup_{B \in \mathcal{B}_2} E(B)$, which is m_2 -covered by \mathcal{T} , or E_∂ , which is m_3 -covered by \mathcal{T} . Therefore, $E(G)$ can be covered by at most $m_1 + m_2 + m_3 + n$ sets in \mathcal{T} , contradicting the assumption that \mathcal{T} has covering number at least $m_1 + m_2 + m_3 + n + 1$. □

4.6.2 From disjoint connected subgraphs to an initial necklace

The following fact is just a slight refinement of the simple observation that a large connected graph must have a long path or a vertex of high degree.

Lemma 4.6.3. *There exists a function $f_{4.6.3} : \mathbb{N} \rightarrow \mathbb{N}$ such that, for $m, n \in \mathbb{N}$ with $m = f_{4.6.3}(n)$, if G is a connected graph, and $\{H_1, \dots, H_m\}$ are vertex-disjoint, connected, non-empty subgraphs of G , then G contains vertex-disjoint, connected subgraphs $\{H'_1, \dots, H'_n\}$ supported by $\{H_1, \dots, H_m\}$ such that either*

(i) G contains a hub z of $\{H'_1, \dots, H'_n\}$ or

(ii) for each $i \in \{1, \dots, n-1\}$, G contains an edge between H'_i and H'_{i+1} .

Proof. Let

$$f_{4.6.3}(n) = \frac{1 - (n-2)^{n-2}}{3-n}$$

For each $i \in \{1, \dots, m\}$, let T_i be a spanning tree of H_i . Because $\{T_1, \dots, T_m\}$ are vertex disjoint, there exists a spanning tree T of G such that, for each $i \in \{1, \dots, n\}$, T_i is a subtree of T .

Let T' be obtained from T by contracting, for each $i \in \{1, \dots, m\}$, $E(T_i)$ to obtain a single vertex v_i . Let \tilde{T} be the smallest subtree of T' containing $\{v_1, \dots, v_m\}$. Note that it suffices to find a family $\{\tilde{T}_1, \dots, \tilde{T}_n\}$ of disjoint subtrees of \tilde{T} such that, for each $i \in \{1, \dots, n\}$, $V(\tilde{T}_i) \cap \{v_1, \dots, v_m\} \neq \emptyset$ and either \tilde{T} contains a hub z of $\{\tilde{T}_1, \dots, \tilde{T}_n\}$ or, for each $i \in \{1, \dots, n-1\}$, \tilde{T} contains an edge between \tilde{T}_i and \tilde{T}_{i+1} .

Note that, by minimality of \tilde{T} , each leaf of \tilde{T} is in $\{v_1, \dots, v_m\}$. If \tilde{T} has a vertex z of degree at least n , then let $\{\tilde{T}_1, \dots, \tilde{T}_n\}$ be distinct, non-empty components of $\tilde{T} - z$; z is a hub of $\{\tilde{T}_1, \dots, \tilde{T}_n\}$ and for each $i \in \{1, \dots, n\}$, \tilde{T}_i contains some leaf of \tilde{T} , which is in $\{v_1, \dots, v_m\}$. Thus, we may assume that each vertex in \tilde{T} has degree at most $n-1$.

For $u, v \in V(\tilde{T})$, let $P_{u,v}$ be the path in \tilde{T} from u to v and let $\tilde{d}(u, v)$ be the number of vertices in $P_{u,v}$ that do not have degree 2 in \tilde{T} . Let x be a leaf of \tilde{T} . For each $i \geq 2$, the number of vertices $v \in V(\tilde{T})$ with $\tilde{d}(x, v) = i$ is at most $(n-2)^{i-2}$, so the number of vertices $v \in V(\tilde{T})$ with $\tilde{d}(x, v) \leq n-1$ is at most

$$1 + \sum_{i=2}^{n-1} (n-2)^{i-2} = 1 + \frac{1 - (n-2)^{n-2}}{3-n} > m.$$

Therefore \tilde{T} contains a path Q with at least n vertices, y_1, \dots, y_n that do not have degree 2; we may assume that y_1, \dots, y_n occur in that order along Q . For each $i \in \{1, \dots, n-1\}$, let Q_i be the subpath of Q between y_i and y_{i+1} .

For each $i \in \{1, \dots, n\}$, there exists a path R_i from y_i to a leaf such that R_i is internally vertex disjoint from Q —either y_i is a leaf, in which case $y_i \in \{v_1, \dots, v_m\}$, or y_i has degree at least 3, in which case there is a subtree attached to y_i that contains a leaf and is disjoint from Q . For each $i \in \{1, \dots, n-1\}$, let $\tilde{T}_i = R_i \cup (Q_i - \{y_{i+1}\})$ and let $\tilde{T}_n = R_n$. Then $\{\tilde{T}_1, \dots, \tilde{T}_n\}$ is a collection of vertex-disjoint subtrees of \tilde{T} each containing a vertex in $\{v_1, \dots, v_m\}$, and, for each $i \in \{1, \dots, n-1\}$, \tilde{T} contains an edge from \tilde{T}_i to \tilde{T}_{i+1} , as desired. \square

The basic idea now is that alternative (ii) in Lemma 4.6.3 will give us a $(1, 0, n)$ -necklace where each bead contains at least θ vertices of our θ -connected set. Alternative (i), on the other hand, gives us a hub, and if we are able to find θ hubs, then we will have a $K_{\theta, n}$ -minor. Proposition 4.6.4 handles this basic induction argument.

Proposition 4.6.4. *There exists a function $f_{4.6.4} : \mathbb{N}^3 \rightarrow \mathbb{N}$ such that, for $\theta, \ell, m, n \in \mathbb{N}$ with $\theta \geq 3$, $n \geq 3$ and $m = f_{4.6.4}(\theta, \ell, n)$, if G is a graph, \mathcal{T} is a θ -tangle in G , \mathcal{H} is a family of disjoint, connected subgraphs of G , none of which are 1-covered by \mathcal{T} , Z is a set of ℓ hubs of \mathcal{H} , then G contains either*

- (i) *a family \mathcal{H}' of disjoint, connected subgraphs supported by \mathcal{H} with $|\mathcal{H}'| \geq n$ and a set Z' of hubs of \mathcal{H}' with $|Z'| \geq \theta$ or*
- (ii) *a $(1, 0, n)$ -necklace $(B_i : i \in \mathbb{Z}_n)$ supported by \mathcal{H} .*

Proof. By induction on $\theta - \ell$. If $\ell \geq \theta$, then alternative (i) holds with $\mathcal{H}' = \mathcal{H}$ and $Z' = Z$, so we may assume $\ell < \theta$.

By Lemma 4.6.1, G contains a unique Z -bridge B_0 for which $E(G) - E(B_0) \in \mathcal{T}$. Let G_0 be the connected component of $G - Z$ contained in B_0 .

Claim 4.6.4.1. *For each $H \in \mathcal{H}$, H is a subgraph of G_0*

Proof of Claim. For each $H \in \mathcal{H}$, $V(H) \cap Z = \emptyset$ and H is connected, so H is contained in a Z -bridge; but H is not 1-covered by \mathcal{T} , so $E(H) \not\subseteq E(G) - E(B_0)$, so $E(H) \subseteq E(B_0)$. Note that H does not contain any edge incident with a vertex in Z , so H is a subgraph of G_0 , proving the claim. \square (Claim)

Thus, G_0 is a connected graph and \mathcal{H} is a collection of at least $m = f_{4.6.3}(n')$ connected subgraphs of G_0 , so, by Lemma 4.6.3, G_0 contains a family of disjoint, connected subgraphs $\mathcal{H}' = \{H'_1, \dots, H'_{n'}\}$ supported by \mathcal{H} such that either

- (a) G_0 contains a hub z of \mathcal{H}' or
- (b) for each $i \in \{1, \dots, n' - 1\}$, G_0 contains an edge between H'_i and H'_{i+1} .

If case (a) holds, then $Z \cup \{z\}$ is a collection of $\ell + 1$ hubs of \mathcal{H}' . Because \mathcal{H}' is supported by \mathcal{H} , for each $i \in \{1, \dots, n'\}$, H'_i contains a set that is not 1-covered by \mathcal{T} , so H'_i is not 1-covered by \mathcal{T} . Because $n' \geq f_{4.6.4}(\theta, \ell + 1, n)$, applying the induction hypothesis to \mathcal{H}' shows that the proposition holds.

Now suppose case (b) holds, and note that $n' \geq n$. For $i \in \{1, \dots, n-1\}$, let e_i be an edge between H'_i and H'_{i+1} , let $u_i \in V(H'_i)$ and $v_{i+1} \in V(H'_{i+1})$ be incident with e_i , and let $B_i = H'_i \cup \{e_i\}$. Let $v_1 \in V(H'_1)$ and let $B_0 = H'_n \cup \{v_1\}$.

Claim 4.6.4.2. $(B_i : i \in \mathbb{Z}_n)$ is a $(1, 0, n)$ -necklace.

Proof of Claim. For each $i \in \mathbb{Z}_n$, $V(B_i) \cap V(B_{i+1}) = \{v_{i+1}\}$, so necklace axiom (N1) holds. Note that $v_1 \in V(B_1) - V(B_2)$, so $V(B_0) \cap V(B_2) = \emptyset$; the other pairs of non-adjacent beads are disjoint because the graphs in \mathcal{H}' are disjoint, so axiom (N2) holds. For each $i \in \{1, \dots, n-1\}$, B_i is connected, and hence contains a path between v_i and v_{i+1} , so $(B_i, \{v_i\}, \{v_{i+1}\})$ is a 1-bead, so axiom (N3) holds. Finally, note that B_0 is the union of two connected subgraphs: $\{v_1\}$ and H'_n ; because $v_n \in H'_n$, $(B_0, \{v_1\}, \{v_n\})$ is a 0-bead, so axiom (N4) holds. \square (Claim)

Thus, $(B_i : i \in \mathbb{Z}_n)$ is a $(1, 0, n)$ -necklace supported by \mathcal{H}' , and hence supported by \mathcal{H} , proving the proposition. \square

From Proposition 4.6.4 we are now able to find a $K_{\theta, n}$ -minor or a $(1, 0, n)$ -necklace with θ paths between each bead:

Lemma 4.6.5. *There exists a function $f_{4.6.5} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that, for $\theta, m, n \in \mathbb{N}$ with $n \geq \theta \geq 3$ and $m = f_{4.6.5}(\theta, n)$, if G is a graph and \mathcal{T} is a θ -tangle in G with covering number at least m , then either G contains a \mathcal{T} -aligned $K_{\theta, n}$ -minor or G contains a \mathcal{T} -aligned $(1, 0, n)$ -necklace $(B_i : i \in \mathbb{Z}_n)$ such that for each $i, j \in \{1, \dots, n\}$, G contains θ vertex-disjoint $(V(B_i), V(B_j))$ -paths.*

Proof. Define $f_{4.6.5}(\theta, n) = f_{4.6.2}(\theta, f_{4.6.4}(\theta, 0, n))$. Let $\theta, n \in \mathbb{N}$, $m = f_{4.6.5}(\theta, n)$. Suppose G is a graph with a θ -tangle \mathcal{T} and covering number at least m .

By Lemma 4.6.2, there exists a collection \mathcal{H} of disjoint, connected subgraphs of G , none of which are 1-covered by \mathcal{T} such that $|\mathcal{H}| \geq f_{4.6.4}(\theta, 0, n)$. By Proposition 4.6.4, G contains either

- (a) a family \mathcal{H}' of disjoint, connected subgraphs supported by \mathcal{H} with $|\mathcal{H}'| \geq n$ and a set Z' of hubs of \mathcal{H}' with $|Z'| \geq \theta$ or
- (b) a $(1, 0, n)$ -necklace $(B_i : i \in \mathbb{Z}_n)$ supported by \mathcal{H} .

Claim 4.6.5.1. *If case (a) holds, then G contains a \mathcal{T} -aligned $K_{\theta, n}$ -minor.*

Proof of Claim. For $H \in \mathcal{H}'$ and $z \in Z'$, choose an edge $e_{z,H}$ between z and H . Contracting each $H \in \mathcal{H}$ to a single vertex v_H and deleting all other edges except $\{e_{z,H} : z \in Z', H \in \mathcal{H}'\}$ gives a minor G' isomorphic to $K_{\theta,n}$; we need only show that G' is \mathcal{T} -aligned. Fix a $(\theta - 1)$ -separation (X, Y) in G . It cannot be the case that $Z' \subseteq \partial(X)$, so, without loss of generality, there is some $z_0 \in Z'$ such that each edge incident with z_0 is in Y . In particular, $e_{z_0,H} \in Y$ for each $H \in \mathcal{H}'$, so H contains a vertex incident with an edge of Y . If H also contains a vertex incident with an edge of X then, because H is connected, H contains some vertex in $\partial(X)$. Because the subgraphs in \mathcal{H}' are vertex-disjoint, at most $\theta - 1$ of them can contain a vertex in $\partial(X)$, so at most $\theta - 1$ of them contain a vertex incident with an edge in X .

Because $|\mathcal{H}'| = n \geq \theta$, there is some $H \in \mathcal{H}'$ such that $V(H) \cap V(X) = \emptyset$, so $E(H) \subseteq Y$. Because \mathcal{H}' is supported by \mathcal{H} , and no subgraph in \mathcal{H} is 1-covered by \mathcal{T} , H is also not 1-covered by \mathcal{T} , so $Y \notin \mathcal{T}$. Therefore, $X \in \mathcal{T}$.

There are at most $\theta - 1$ vertices $z \in Z'$ such that $\{e_{z,H} : H \in \mathcal{H}'\} \cap X \neq \emptyset$, and, for each of these, $e_{z_0,H} \in Y$, and

$$|\{e_{z,H} : H \in \mathcal{H}'\} \cap X| \leq \theta - 1,$$

therefore,

$$|E(G') \cap X| = |\{e_{z,H} : z \in Z', H \in \mathcal{H}'\} \cap X| \leq (\theta - 1)^2.$$

On the other hand,

$$|E(G')| \geq |\mathcal{H}'||Z'| = n\theta \geq 2(\theta - 1)\theta > 2(\theta - 1)^2 \geq 2|E(G') \cap X|,$$

so X , the side of the $(\theta - 1)$ -separation in \mathcal{T} , contains fewer than half of the edges in G' . Hence, G' is \mathcal{T} -aligned, proving the claim. \square (Claim)

We may now assume (b) holds, so G contains a $(1, 0, n)$ -necklace $(B_i : i \in \mathbb{Z}_n)$ supported by \mathcal{H} .

Claim 4.6.5.2. *For each $i, j \in \{1, \dots, n\}$, G contains θ vertex-disjoint paths between B_i and B_j .*

Proof of Claim. Suppose not, so G has a $(\theta - 1)$ -separation (X, Y) with $E(B_i) \subseteq X$ and $E(B_j) \subseteq Y$. One of X or Y must be in \mathcal{T} by tangle axiom (T2); without loss of generality, $X \in \mathcal{T}$. But $(B_i : i \in \mathbb{Z}_n)$ is supported by \mathcal{H} , so there is some $H \in \mathcal{H}$ such that $H \subseteq B_i \subseteq X \in \mathcal{T}$, contradicting the fact that H is not 1-covered by \mathcal{T} . \square (Claim)

Therefore, $(B_i : i \in \mathbb{Z}_n)$ is a \mathcal{T} -aligned $(1, 0, n)$ -necklace in G with θ vertex-disjoint paths between any two beads, so the lemma holds in this case as well. \square

4.7 Jumps in a necklace

Now that we have a necklace with θ disjoint paths between each bead, we need to understand how those paths can interact with the necklace. A path between two beads might go through several other beads along the way, but it will be composed of some subpaths within beads together with *jumps* between the beads—if $(B_i : i \in \mathbb{Z}_n)$ is a (t, s, n) -necklace in a graph G and $i, j \in \mathbb{Z}_n$, then an (i, j) -jump of $(B_i : i \in \mathbb{Z}_n)$ in G is a path P in G with endpoints x and y such that $x \in V(B_i)$, $y \in V(B_j)$ and $V(P) \cap \bigcup_{i'=1}^n V(B_{i'}) = \{x, y\}$. A necklace $(B_i : i \in \mathbb{Z}_n)$ is said to be *jump-free* in G if, for each $i, j \in \mathbb{Z}_n$ with $j - i \notin \{-1, 0, 1\}$, G does not contain any (i, j) -jump of $(B_i : i \in \mathbb{Z}_n)$.

Lemma 4.7.1 shows that we only need to consider two possible configurations of jumps of our necklace: we will be able to assume that the necklace is either jump-free or that there is a jump from the weak bead to each other bead in the necklace.

Lemma 4.7.1. *There exists a function $f_{4.7.1} : \mathbb{N} \rightarrow \mathbb{N}$ such that, for $t, s, m, n \in \mathbb{N}$ with $t \geq \max\{s, 1\}$, $n \geq 3$ and $m = f_{4.7.1}(n)$, if G is a graph containing a (t, s, m) -necklace $(B_i : i \in \mathbb{Z}_m)$ then either*

- (i) G contains a (t, s, n) -necklace $(B'_i : i \in \mathbb{Z}_n)$ supported by $(B_i : i \in \mathbb{Z}_m)$ such that $(B'_i : i \in \mathbb{Z}_n)$ is jump-free in G or
- (ii) G contains a (t, s, n) -necklace $(B'_i : i \in \mathbb{Z}_n)$ supported by $(B_i : i \in \mathbb{Z}_m)$ such that for each $i \in \{1, \dots, n-1\}$, G contains a $(0, i)$ -jump.

Proof. Let $\tilde{n} = (n-1)^2 + 1$ and let $f_{4.7.1}(n) = (\tilde{n}-1)^{n-2} + 2$. Choose a sequence of natural numbers $(a_1, \dots, a_\ell, b_\ell, \dots, b_1)$ such that

$$1 \leq a_1 < a_2 < \dots < a_\ell < b_\ell < b_{\ell-1} < \dots < b_1 \leq m$$

and, for each $i \in \{1, \dots, \ell\}$, $b_i - a_i \geq (\tilde{n}-1)^{n-i-1} + 1$ and G contains an (a_i, b_i) -jump of $(B_i : i \in \mathbb{Z}_m)$, and ℓ is maximized.

Suppose first that $\ell \geq n-1$. Let $a_0 = 0$ and, for each $i \in \{1, \dots, n-1\}$, let $B'_i = \bigcup_{j=a_{i-1}+1}^{a_i} B_j$ and let $B'_0 = \bigcup_{j=a_{n-1}+1}^m B_j$. By Lemma 4.5.6, $(B'_i : i \in \mathbb{Z}_n)$ is a (t, s, n) -necklace supported by $(B_i : i \in \mathbb{Z}_m)$ and, for each $i \in \{1, \dots, n-1\}$, G contains an $(0, i)$ -jump of $(B'_i : i \in \mathbb{Z}_n)$, so alternative (ii) holds.

Therefore, we may assume $\ell < n - 1$. Because $t \geq 1$, G contains a $(0, 1)$ -jump of $(B_i : i \in \mathbb{Z}_m)$ (namely, the path of length 0 at a vertex of $V(B_0) \cap V(B_1)$), and $m - 1 \geq (\tilde{n} - 1)^{n-2} + 1$, so $\ell \geq 1$. Then $b_\ell - a_\ell \geq (\tilde{n} - 1)^{n-\ell-1} + 1$ and for each $a', b' \in \{a_\ell + 1, \dots, b_\ell - 1\}$ with $b' - a' \geq (\tilde{n} - 1)^{n-\ell-2} + 1$, G does not contain an (a', b') -jump.

Let $\alpha = (\tilde{n} - 1)^{n-\ell-2} + 1$. For $i \in \{1, \dots, \tilde{n} - 1\}$, let

$$\tilde{B}_i = \bigcup_{j=1}^{\alpha-1} B_{a_\ell + (i-1)(\alpha-1) + j}$$

and let

$$\tilde{B}_0 = \bigcup_{j=a_\ell + (\tilde{n}-1)(\alpha-1) + 1}^{m-1} B_j \cup \bigcup_{j=0}^{a_\ell} B_j,$$

so $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$ is a (t, s, \tilde{n}) -necklace by Lemma 4.5.6. For each $\tilde{a}, \tilde{b} \in \{1, \dots, \tilde{n} - 1\}$ with $\tilde{b} - \tilde{a} \geq 2$, any (\tilde{a}, \tilde{b}) -jump of $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$ is an (a, b) -jump of $(B_i : i \in \mathbb{Z}_m)$ for some

$$a \in \{a_\ell + (\tilde{a} - 1)(\alpha - 1) + 1, \dots, a_\ell + \tilde{a}(\alpha - 1)\}$$

and some

$$b \in \{a_\ell + (\tilde{b} - 1)(\alpha - 1) + 1, \dots, a_\ell + \tilde{b}(\alpha - 1)\},$$

but $b - a \geq \alpha$ and

$$a_\ell + 1 \leq a < b \leq a_\ell + (\tilde{n} - 1)(\alpha - 1) = a_\ell + (\tilde{n} - 1)^{n-\ell-1} \leq b_\ell - 1,$$

contradiction. Therefore, for each $\tilde{a}, \tilde{b} \in \{1, \dots, \tilde{n} - 1\}$ with $\tilde{b} - \tilde{a} \geq 2$, G does not contain a (\tilde{a}, \tilde{b}) -jump of $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$.

Suppose there exists some $a \in \{1, \dots, \tilde{n} - n - 2\}$ such that for each $i \in \{a + 1, \dots, a + n - 1\}$, G does not contain an $(0, i)$ -jump of $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$. For each $i \in \{1, \dots, n - 1\}$, let $B'_i = \tilde{B}_{a+i}$ and let $B'_0 = \bigcup_{i=a+n}^{\tilde{n}} \tilde{B}_i \cup \bigcup_{i=1}^a \tilde{B}_i$. Then $(B'_i : i \in \mathbb{Z}_n)$ is a (t, s, n) -necklace by Lemma 4.5.6, and is jump-free by construction, so alternative (i) holds.

Finally suppose that, for each $a \in \{1, \dots, \tilde{n} - n - 2\}$ there is some $i \in \{a + 1, \dots, a + n - 1\}$ such that G contains an (i, \tilde{n}) -jump of $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$. For each $i \in \{1, \dots, n - 1\}$, let $B'_i = \bigcup_{j=(i-1)(n-1)+1}^{i(n-1)} \tilde{B}_j$ and let $B'_0 = \tilde{B}_0 = \tilde{B}_{(n-1)^2+1}$. Then $(B'_i : i \in \mathbb{Z}_n)$ is a (t, s, n) -necklace by Lemma 4.5.6, and, for each $i \in \{1, \dots, n - 1\}$, G contains an $(0, i)$ -jump of $(B'_i : i \in \mathbb{Z}_n)$, so alternative (ii) holds. \square

4.7.1 Jump-free necklaces

Jump-free necklaces are convenient because the θ paths between a pair of beads are quite restricted.

We need a small lemma about connectivity, Lemma 4.7.2, which solves the following problem: Suppose we have one family of n disjoint paths that start off in some set X , and another family of n disjoint paths that end up at a set Y , and we want to find a family of paths from X to Y . If we also have a family of n disjoint connected subgraphs that each meet all of these paths, then we can start off on the paths from X , and use the connected subgraphs to switch to the paths to Y . This ability to switch from one family of paths to another turns out to be quite useful. In Lemma 4.7.2, you should think of the family \mathcal{A} as the paths that start in X , \mathcal{B} as the paths that end up in Y , and \mathcal{H} as the graphs used to switch between the paths.

Lemma 4.7.2. *If $n \in \mathbb{N}$ and $\mathcal{A}, \mathcal{B}, \mathcal{H}$ are three families of connected pairwise vertex-disjoint subgraphs of a graph G such that $\min\{|\mathcal{A}|, |\mathcal{B}|, |\mathcal{H}|\} \geq n$ and for each $A \in \mathcal{A}$, $B \in \mathcal{B}$, $H \in \mathcal{H}$, $V(A) \cap V(H) \neq \emptyset$ and $V(B) \cap V(H) \neq \emptyset$ and if $X, Y \subseteq V(G)$ such that, for each $A \in \mathcal{A}$, $X \cap V(A) \neq \emptyset$ and for each $B \in \mathcal{B}$, $Y \cap V(B) \neq \emptyset$, then G contains n vertex-disjoint (X, Y) -paths.*

Proof. Suppose G does not contain n vertex-disjoint (X, Y) -paths, and let (G_1, G_2) be an $(n - 1)$ -separation in G such that $X \subseteq V(G_1)$ and $Y \subseteq V(G_2)$. Note that, for each $A \in \mathcal{A}$,

$$V(A) \cap V(G_1) \supseteq V(A) \cap X \neq \emptyset,$$

so, if $V(A) \cap V(G_2) \neq \emptyset$ as well then, because A is connected, $V(A)$ contains a vertex in $V(G_1) \cap V(G_2)$; also,

$$|V(G_1) \cap V(G_2)| < n \leq |\mathcal{A}|,$$

so there exists some $A_0 \in \mathcal{A}$ such that $V(A_0) \cap V(G_2) = \emptyset$. Similarly, there exists some $B_0 \in \mathcal{B}$ such that $V(B_0) \cap V(G_1) = \emptyset$. For each $H \in \mathcal{H}$,

$$V(H) \cap V(G_1) \supseteq V(H) \cap V(A_0) \neq \emptyset$$

and

$$V(H) \cap V(G_2) \supseteq V(H) \cap V(B_0) \neq \emptyset,$$

so H contains some vertex in $V(G_1) \cap V(G_2)$. But $|\mathcal{H}| \geq n$, contradicting the fact that (G_1, G_2) is an $(n - 1)$ -separation. \square

Before looking at jump-free necklaces, we first need to understand how surplus paths can be used to augment a jump-free “piece” of a necklace. If (B_0, \dots, B_{n+1}) is a sequence of subgraphs of a graph G , for each $i \in \{0, \dots, n\}$, $|V(B_i) \cap V(B_{i+1})| = t$, for each $i, j \in \{0, \dots, n+1\}$ with $i < j - 1$, $V(B_i) \cap V(B_j) = \emptyset$ and, for each $i \in \{1, \dots, n\}$, $(B_i, V(B_{i-1}) \cap V(B_i), V(B_i) \cap V(B_{i+1}))$ is a t -bead, then (B_0, \dots, B_{n+1}) is a (t, n) -chain.

Lemma 4.7.3. *There exists a function $f_{4.7.3} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that, for $t, m, n \in \mathbb{N}$ with $t \geq 1$, $n \geq 3$, and $m = f_{4.7.3}(t, n)$, if G is a graph, (B_0, \dots, B_{m+1}) is a jump-free (t, m) -chain in G , and G contains a collection of $t + 1$ vertex-disjoint (B_1, B_m) -paths, then G contains a $(t + 1, n)$ -chain (B'_0, \dots, B'_{n+1}) supported by (B_0, \dots, B_{m+1}) . Moreover,*

(i) for each $i \in \{1, \dots, n\}$, $V(B'_i) \cap V(B_0) = V(B'_i) \cap V(B_{m+1}) = \emptyset$,

(ii) $B'_0 \cap B_0 = B_1 \cap B_0$ and $(B'_0, B'_0 \cap B_0, B'_0 \cap B_1)$ is a t -bead, and

(iii) $B'_{n+1} \cap B_{m+1} = B_m \cap B_{m+1}$ and $(B'_{n+1}, B'_{n+1} \cap B_{m+1}, B'_{n+1} \cap B'_n)$ is a t -bead.

Proof. Let $f_{4.7.3}(t, n) = 4t(n + 1) + 2t + 1$. The major difficulty in this proof is that the $t + 1$ vertex-disjoint (B_1, B_m) -paths might wind back and forth through the other beads in complex ways, so we need to choose these paths very carefully. Let \mathcal{P} be a collection of $t + 1$ vertex-disjoint (B_2, B_{m-1}) paths minimizing

$$\left| \bigcup_{P \in \mathcal{P}} E(P) - \bigcup_{i=2}^{(m+1)/2} E(B_{2i}) \right|.$$

Notice first that, because (B_0, \dots, B_{m+1}) is jump-free, for each $i \in \{2, \dots, m-1\}$, $G - V(B_i)$ has no (B_2, B_{m-1}) -path, so for each $P \in \mathcal{P}$, $V(P) \cap V(B_i) \neq \emptyset$.

For each $P \in \mathcal{P}$, let x_P be the endpoint of P in B_2 and let y_P be the endpoint of P in B_{m-1} .

For this proof, if P is a path and $x, y \in V(P)$, we use the notation $P[x, y]$ to denote the subpath from x to y in P .

Now we need to show that the paths in \mathcal{P} cannot backtrack too far through the beads.

Claim 4.7.3.1. *If $P \in \mathcal{P}$, $i, j \in \{1, \dots, (m - 1)/2\}$, $i + t \leq j$, $v \in V(P) \cap V(B_{2j})$, then $P[v, y_P]$ is disjoint from B_{2i} .*




Figure 4.11: Illustration of notation used in the proof of Lemma 4.7.3

Proof of Claim. Suppose the contrary and let $u \in V(P[v, y_P]) \cap V(B_{2i})$. Let

$$G' = \bigcup_{i'=i}^j B_{2i'} \cup \bigcup ((\mathcal{P} - \{P\}) \cup \{P[x_P, v], P[u, y_P]\}).$$

See Figure 4.11. By Lemma 4.7.2 with $\mathcal{A} = (\mathcal{P} - \{P\}) \cup \{P[x_P, v]\}$, $\mathcal{B} = (\mathcal{P} - \{P\}) \cup \{P[u, y_P]\}$ and $\mathcal{H} = \{B_{2i'} : i' \in \{i, \dots, j\}\}$, G' contains a collection \mathcal{P}' of $t + 1$ vertex-disjoint $(\{x_{P'} : P' \in \mathcal{P}'\}, \{y_{P'} : P' \in \mathcal{P}'\})$ -paths. But

$$\left| \bigcup_{P' \in \mathcal{P}'} E(P') - \bigcup_{i'=1}^{(m-1)/2} E(B_{2i'}) \right| \leq \left| \bigcup_{P' \in \mathcal{P}} E(P') - \bigcup_{i'=1}^{(m-1)/2} E(B_{2i'}) \right| - \left| E(P[u, v]) - \bigcup_{i'=1}^{(m-1)/2} E(B_{2i'}) \right|$$

and B_{2i} and B_{2j} are vertex-disjoint so $E(P[u, v])$ must have an edge outside of $\bigcup_{i'=1}^{(m-1)/2} E(B_{2i'})$, contradicting minimality of \mathcal{P} . Hence we have proven the claim that if $P \in \mathcal{P}$, $i, j \in \{1, \dots, (m-1)/2\}$, $i + t \leq j$, $v \in V(P) \cap V(B_{2j})$, then $P[v, y_P]$ is disjoint from B_{2i} . \square (Claim)

For $i \in \{2, \dots, m-1\}$, let $x_P^i \in V(P) \cap V(B_i)$ minimizing the distance along P from x_P to x_P^i . Therefore, for each $i, j \in \{1, \dots, (m-1)/2\}$ with $i + t \leq j$, B_{2i} is vertex disjoint from $P[x_P^{2j}, y_P]$.

For $i \in \{1, \dots, n\}$, let

$$B'_i = \bigcup_{j=1}^t B_{4ti+2j} \cup \bigcup_{P \in \mathcal{P}} P[x_P^{4ti+2}, x_P^{4t(i+1)+2}]$$

so, for each $i \in \{1, \dots, n-1\}$, $B'_i \cap B'_{i+1} = \{x_P^{4t(i+1)+2} : P \in \mathcal{P}\}$, which has size $t+1$.

Notice that, for $i, i' \in \{1, \dots, n\}$ if $i < i'$ or $i > i'+1$ then $\bigcup_{j=1}^t B_{4ti+2j}$ is disjoint from $\bigcup_{P \in \mathcal{P}} P[x_P^{4ti+2}, x_P^{4t(i'+1)+2}]$. Therefore, if $|i - i'| > 1$, then $V(B'_i) \cap V(B'_{i'}) = \emptyset$.

Claim 4.7.3.2. For each $i \in \{1, \dots, n\}$,

$$(B'_i, \{x_P^{4ti+2} : P \in \mathcal{P}\}, \{x_P^{4t(i+1)+2} : P \in \mathcal{P}\})$$

is a $(t+1)$ -bead.

Proof of Claim. Let $X \subseteq \{x_P^{4ti+2} : P \in \mathcal{P}\}$ and $Y \subseteq \{x_P^{4t(i+1)+2} : P \in \mathcal{P}\}$ such that $|X| = |Y| \leq t+1$. If $|X| = t+1$ then $X = \{x_P^{4ti+2} : P \in \mathcal{P}\}$ and $Y = \{x_P^{4t(i+1)+2} : P \in \mathcal{P}\}$, so $\{P[x_P^{4ti+2}, x_P^{4t(i+1)+2}] : P \in \mathcal{P}\}$ is a family of $|X|$ vertex-disjoint (X, Y) -paths in B'_i . Otherwise, $|X| \leq t$, so by Lemma 4.7.2 with $\mathcal{A} = \{P[x_P^{4ti+2}, x_P^{4t(i+1)+2}] : x_P^{4ti+2} \in X\}$, $\mathcal{B} = \{P[x_P^{4ti+2}, x_P^{4t(i+1)+2}] : x_P^{4ti+2} \in X\}$ and $\mathcal{H} = \{B_{4ti+2j} : j \in \{1, \dots, t\}\}$, B'_i contains $|X|$ vertex-disjoint (X, Y) -paths. Hence $(B'_i, \{x_P^{4ti+2} : P \in \mathcal{P}\}, \{x_P^{4t(i+1)+2} : P \in \mathcal{P}\})$ is a $(t+1)$ -bead. \square (Claim)

Claim 4.7.3.3. (B'_0, \dots, B'_{n+1}) is a $(t+1, n)$ -chain.

Proof of Claim. Notice that, by minimality of $|\bigcup_{P \in \mathcal{P}} E(P) - \bigcup_{i=2}^{(m+1)/2} E(B_{2i})|$, for each $P \in \mathcal{P}$, P intersects B_2 only at x_P , so P is disjoint from B_0 , and P intersects B_{m-1} only at y_P , so P is disjoint from B_{m+1} . Therefore, for each $i \in \{1, \dots, n\}$, $B'_i \cap B_0 = B'_i \cap B_{m+1} = \emptyset$, so (i) holds. Let

$$B'_0 = \bigcup_{j=1}^{2t} B_j \cup \bigcup_{P \in \mathcal{P}} P[x_P, x_P^{4t+2}]$$

and

$$B'_{n+1} = \bigcup_{j=4t(n+1)+1}^m B_j \cup \bigcup_{P \in \mathcal{P}} P[x_P^{4t(n+1)+2}, y_P].$$

Then $V(B'_0) \cap V(B'_1) = \{x_P^{4t+2} : P \in \mathcal{P}\}$, $V(B'_{n+1}) \cap V(B'_n) = \{x_P^{4t(n+1)+2} : P \in \mathcal{P}\}$, for each $i \in \{2, \dots, n+1\}$ $V(B'_0) \cap V(B'_i) = \emptyset$ and for each $i \in \{0, \dots, n-1\}$, $V(B'_{n+1}) \cap V(B'_i) = \emptyset$, so (B'_0, \dots, B'_{n+1}) is a $(t+1, n)$ -chain. \square (Claim)

Claim 4.7.3.4. $(B'_0, V(B'_0) \cap V(B_0), V(B'_0) \cap V(B'_1))$ is a t -bead.

Proof of Claim. Let $X \subseteq V(B'_0) \cap V(B_0) = V(B_0) \cap V(B_1)$ and $Y \subseteq V(B'_0) \cap V(B'_1) = \{x_P^{4t+2} : P \in \mathcal{P}\}$. Choose a collection \mathcal{Q} of t vertex-disjoint $(V(B_0) \cap V(B_1), V(B_{2t}))$ -paths in $\bigcup_{j=1}^{2t} B_j$. Because $|V(B_0) \cap V(B_1)| = t = |\mathcal{Q}|$, each vertex in $V(B_0) \cap V(B_1)$ is in exactly one path in \mathcal{Q} . By Lemma 4.7.2 with $\mathcal{A} = \{Q : Q \cap X \neq \emptyset\}$, $\mathcal{B} = \{P[x_P, x_P^{4t+2}] : x_P^{4t+2} \in Y\}$ and $\mathcal{H} = \{B_{2j} : j \in \{1, \dots, t\}\}$, B'_0 contains t vertex-disjoint (X, Y) -paths, and hence $(B'_0, V(B'_0) \cap V(B_0), V(B'_0) \cap V(B'_1))$ is a t -bead. \square (Claim)

Therefore, (ii) holds.

Finally $V(B'_{n+1}) \cap V(B_{m+1}) = V(B_m) \cap V(B_{m+1})$ and, because $m - (4t(n+1) + 1) = 2t$, an argument similar to Claim 4.7.3.4 shows that $(B'_0, V(B'_n) \cap V(B'_{n+1}), V(B'_{n+1}) \cap V(B_{m+1}))$ is a t -bead, so (iii) holds. \square

Now we are able to show that surplus paths between the beads of a jump-free necklace can be used to increase its connectivity.

Lemma 4.7.4. *There exists a function $f_{4.7.4} : \mathbb{N}^3 \rightarrow \mathbb{N}$ such that, for $t, s, m, n \in \mathbb{N}$ with $t \geq \max\{s, 1\}$, $n \geq 3$ and $m = f_{4.7.4}(t, s, n)$, if G is a graph, $(B_i : i \in \mathbb{Z}_m)$ is a jump-free (t, s, m) -necklace in G and, for each $i, j \in \mathbb{Z}_m$, G contains a collection of $(t + s + 1)$ vertex-disjoint $(V(B_i), V(B_j))$ -paths, then there exists $t', s' \in \mathbb{N}$ such that $t' + s' = t + s + 1$ and G contains a (t', s', n) -necklace $(B'_i : i \in \mathbb{Z}_n)$ supported by $(B_i : i \in \mathbb{Z}_m)$.*

Proof. Let $a = \max\{\lceil f_{4.7.3}(t, n - 1)/2 \rceil, 2s + 2\}$ and let $f_{4.7.4}(t, s, n) = 4a$. Let $b = m - a$ and let \mathcal{P} be a collection of $t + s + 1$ vertex-disjoint $(V(B_a), V(B_b))$ -paths in G minimizing $|\bigcup_{P \in \mathcal{P}} E(P)|$. If $P \in \mathcal{P}$, $i \in \{a + 2, \dots, b - 2\}$, $j \in \{1, \dots, a - 2\} \cup \{b + 2, \dots, m\}$, $x \in V(P) \cap V(B_i)$ and $y \in V(P) \cap v(B_j)$, then, because $(B_i : i \in \mathbb{Z}_m)$ is jump-free, the subpath P from x to y must contain a vertex z in $V(B_a)$ or $V(B_b)$. But $z \notin \{x, y\}$, so z is not an endpoint of P , so P contains a proper (B_a, B_b) -subpath, contradicting minimality of $|\bigcup_{P \in \mathcal{P}} E(P)|$. On the other hand, each (B_a, B_b) -path in G must contain a vertex of either B_{a+2} or B_{a-2} . Therefore, \mathcal{P} can be partitioned into $\mathcal{P}_1 = \{P \in \mathcal{P} : V(P) \cap V(B_{a+2}) \neq \emptyset\}$ and $\mathcal{P}_2 = \{P \in \mathcal{P} : V(P) \cap V(B_{a-2}) \neq \emptyset\}$. Note that $|\mathcal{P}_1| + |\mathcal{P}_2| = |\mathcal{P}| = t + s + 1$ so either $|\mathcal{P}_1| \geq t + 1$ or $|\mathcal{P}_2| \geq s + 1$.

Claim 4.7.4.1. *If $|\mathcal{P}_1| \geq t + 1$, then G contains a $(t + 1, s, n)$ -necklace supported by $(B_i : i \in \mathbb{Z}_m)$.*

Proof of Claim. Then $(B_{a-1}, \dots, B_{b+1})$ is a $(t, b - a + 1)$ -chain in $G_1 = \bigcup_{i=a-1}^{b+1} B_i \cup \bigcup \mathcal{P}_1$. Note that $(B_{a-1}, \dots, B_{b+1})$ is jump-free in G_1 —if J were an (i, j) -jump of $(B_{a-1}, \dots, B_{b+1})$ in G_1 with $j > i + 1$ then J is a subpath of some $P \in \mathcal{P}_1$; P is disjoint from $\bigcup_{i'=1}^{a-2} B_{i'} \cup$

$\bigcup_{i'=b+2}^m B_{i'}$, so J would also be a $(i+a+1, j+a+1)$ -jump of $(B_i : i \in \mathbb{Z}_m)$, contradicting the fact that $(B_i : i \in \mathbb{Z}_m)$ is jump-free in G . Because $b-a+1 \geq f_{4.7.3}(t, n-1)$, by Lemma 4.7.3, G_1 contains a $(t+1, n-1)$ -chain (B'_0, \dots, B'_n) supported by $(B_{a-1}, \dots, B_{b+1})$ such that for each $i \in \{1, \dots, n-1\}$,

$$V(B'_i) \cap V(B_{a-1}) = V(B'_i) \cap V(B_{b+1}) = \emptyset,$$

$$V(B_{a-1}) \cap V(B'_0) = V(B_{a-1}) \cap V(B_a),$$

$$V(B'_n) \cap V(B_{b+1}) = V(B_b) \cap V(B_{b+1}),$$

$$(B'_0, V(B'_0) \cap V(B_{a-1}), V(B'_0) \cap V(B'_1))$$

is a t -bead, and

$$(B'_n, V(B'_n) \cap V(B_{b+1}), V(B'_{n-1}) \cap V(B'_n))$$

is a t -bead. Then, by Lemma 4.5.5 if $\tilde{B} = \bigcup_{i=1}^{a-1} B_i \cup \bigcup_{i=b+1}^m B_i \cup B'_0 \cup B'_n$, then $(\tilde{B}, V(B'_0) \cap V(B'_1), V(B'_n) \cap V(B'_{n-1}))$ is an s -bead. Hence, $(\tilde{B}, B'_1, \dots, B'_{n-1})$ is a $(t+1, s, n)$ -necklace in G supported by $(B_i : i \in \mathbb{Z}_m)$, as desired. \square (Claim)

If $|\mathcal{P}_2| \geq s+1$ and $s=t$, then the labels of the beads can be shifted to reduce to the case when $|\mathcal{P}_1| \geq t+1$.

Claim 4.7.4.2. *If $|\mathcal{P}_2| \geq s+1$ and $s < t$, then G contains a $(t, s+1, n)$ -necklace supported by $(B_i : i \in \mathbb{Z}_m)$.*

Proof of Claim. Let $\tilde{B} = \bigcup_{i=1}^{a+1} B_i \cup \bigcup_{i=b-1}^m B_i \cup \bigcup \mathcal{P}_2$. Note that $V(\tilde{B}) \cap \bigcup_{i=a+2}^{b-2} V(B_i) = (V(B_{a+1}) \cap V(B_{a+2})) \cup (V(B_{b-2}) \cap V(B_{b-1}))$. It suffices to show that that $(\tilde{B}, V(B_{a+1}) \cap V(B_{a+2}), V(B_{b-2}) \cap V(B_{b-1}))$ is an $(s+1)$ -bead. Note first that \tilde{B} is connected because $\bigcup_{i=1}^{a+1} B_i$ and $\bigcup_{i=b-1}^m B_i$ are both connected and \mathcal{P}_2 is a non-empty collection of paths between these two connected graphs. Let $X \subseteq V(B_{a+1}) \cap V(B_{a+2})$ and $Y \subseteq V(B_{b-2}) \cap V(B_{b-1})$ such that $|X| = |Y| \leq s+1$. It suffices to show that \tilde{B} contains a collection of $s+1$ vertex-disjoint (X, Y) -paths. Let $X_0 = V(B_a) \cap \bigcup_{P \in \mathcal{P}_2} V(P)$ and $Y_0 = V(B_b) \cap \bigcup_{P \in \mathcal{P}_2} V(P)$.

Note that $|X| \leq s+1 \leq t$ so there exists a collection \mathcal{Q} of $|X|$ vertex-disjoint $(X, V(B_m))$ -paths in $\bigcup_{i=1}^{a+1} B_i$. For each $i \in \{1, \dots, s+1\}$, $2i \leq a$ so $B_{2i} \subseteq \tilde{B}$, and, for each $P \in \mathcal{P}_2$ and each $Q \in \mathcal{Q}$, $V(B_{2i}) \cap V(P) \neq \emptyset$ and $V(B_{2i}) \cap V(Q) \neq \emptyset$. By Lemma 4.7.2 with $\mathcal{A} = \mathcal{Q}$, $\mathcal{B} = \mathcal{P}_2$ and $\mathcal{H} = \{B_{2i} : i \in \{1, \dots, s+1\}\}$, \tilde{B} contains a collection \mathcal{P}' of $|X|$ vertex-disjoint (X, Y_0) -paths

Let \mathcal{R} be a collection of t vertex-disjoint $(Y, V(B_m))$ -paths in $\bigcup_{i=b-1}^{m-1} B_i$. For each $i \in \{1, \dots, s+1\}$, $m-2i \geq b$ so $B_{m-2i} \subseteq \tilde{B}$, and, for each $P \in \mathcal{P}_2$ and each $R \in \mathcal{R}$,




Figure 4.12: An illustration of Proposition 4.7.5. (B_0, \dots, B_5) is a $(4, 4)$ -chain, and there are 4 vertex-disjoint paths between $X_1 \cup \{v_3\}$ and X_5 .

$V(B_{m-2i}) \cap V(P) \neq \emptyset$ and $V(B_{m-2i}) \cap V(Q) \neq \emptyset$. By Lemma 4.7.2 with $\mathcal{A} = \mathcal{R}$, $\mathcal{B} = \mathcal{P}'$ and $\mathcal{H} = \{B_{m-2i} : i \in \{1, \dots, s+1\}\}$, \tilde{B} contains a collection of $|X|$ vertex-disjoint (X, Y) -paths.

Hence $(\tilde{B}, V(B_{a+1}) \cap V(B_{a+2}), V(B_{b-2}) \cap V(B_{b-1}))$ is an $(s+1)$ -bead. Thus, $(\tilde{B}, B_{a+2}, \dots, B_{b-2})$ is an $(t, s+1, b-a-2)$ -necklace supported by $(B_i : i \in \mathbb{Z}_m)$ and $b-a-2 \geq n$ so there is some $(t, s+1, n)$ -necklace in G supported by $(B_i : i \in \mathbb{Z}_m)$, as desired. \square (Claim)

In each case, G has the desired necklace, proving the lemma. \square

4.7.2 Necklaces with jumps to the weak bead

The other alternative in Lemma 4.7.1 is that each bead in a (t, s, n) -necklace $(B_i : i \in \mathbb{Z}_n)$ has a jump to the weak bead, B_0 . In this case we won't be able to directly increase $t+s$ (the parameter in the induction), but we will be able to mutate the necklace to maintain the same value of $t+s$ while decreasing the parameter s ; this will be good enough for the induction needed to prove Lemma 4.5.2.

We need a few technical propositions. Propositions 4.7.5, 4.7.6 and 4.7.7 are used to show that, if a subgraph that attaches to a chain of beads in a particular way, then paths can be routed through that subgraph to create a bead with higher connectivity.

Proposition 4.7.5. *If (B_0, \dots, B_5) is a $(t, 4)$ -chain in a graph G and $v_3 \in B_3$, then, for each $X_1 \subseteq V(B_0) \cap V(B_1)$ and each $X_5 \subseteq V(B_4) \cap V(B_5)$ with $|X_1| + 1 = |X_5| \leq t$, G contains $|X_5|$ vertex-disjoint $(X_1 \cup \{v_3\}, X_5)$ -paths; see Figure 4.12.*

Proof. Let $\tilde{t} = |X_5|$. Choose a collection \mathcal{P}_4 of \tilde{t} vertex-disjoint $(V(B_3) \cap V(B_4), X_5)$ -paths in B_4 minimizing $|\bigcup_{P \in \mathcal{P}_4} E(P)|$; let $X_4 = V(B_3) \cap \bigcup_{P \in \mathcal{P}_4} V(P)$, which has size \tilde{t} by minimality of $|\bigcup_{P \in \mathcal{P}_4} E(P)|$. Choose a collection \mathcal{P}_2 of \tilde{t} vertex-disjoint $(V(B_1) \cap V(B_2), V(B_2) \cap V(B_3))$ -paths in B_2 minimizing $|\bigcup_{P \in \mathcal{P}_2} E(P)|$; let $X_3 = V(B_3) \cap \bigcup_{P \in \mathcal{P}_2} V(P)$, which has




Figure 4.13: An illustration of Proposition 4.7.6. One example $X \subseteq (V(B_0) \cap V(B_1)) \cup \{x_1\}$ and $Y \subseteq V(B_4) \cap V(B_5)$ has been shown, along with a collection of vertex disjoint paths between X and Y . The vertices where H meets one of the beads only have one neighbor in H (but possibly other neighbors in the bead).

size \tilde{t} by minimality of $|\bigcup_{P \in \mathcal{P}_2} E(P)|$. Choose a collection \mathcal{P}_3 of \tilde{t} vertex-disjoint (X_3, X_4) -paths in B_3 , and choose a shortest $(\{v_3\}, \bigcup_{P \in \mathcal{P}_3} V(P))$ -path Q in B_3 . Let $P_0^3 \in \mathcal{P}_3$ be the unique path for which $V(P_0^3) \cap V(Q) \neq \emptyset$. Let $P_0^2 \in \mathcal{P}_2$ be the unique path for which $V(P_0^2) \cap V(P_0^3) \neq \emptyset$, and let $P_0^4 \in \mathcal{P}_4$ be the unique path for which $V(P_0^4) \cap V(P_0^3) \neq \emptyset$. Let $X_2 = V(B_1) \cap \bigcup_{P \in \mathcal{P}_2 - \{P_0^2\}} V(P)$, which has size $\tilde{t} - 1$ by minimality of $|\bigcup_{P \in \mathcal{P}_2} E(P)|$. Let \mathcal{P}_1 be a collection of $\tilde{t} - 1$ vertex-disjoint (X_1, X_2) -paths in B_1 .

Then $\bigcup \mathcal{P}_1 \cup \bigcup (\mathcal{P}_2 - \{P_0^2\}) \cup \bigcup (\mathcal{P}_3 - \{P_0^3\}) \cup \bigcup (\mathcal{P}_4 - \{P_0^4\})$ is the union of $\tilde{t} - 1$ vertex-disjoint (X_1, X_5) -paths, and $Q \cup P_0^3 \cup P_0^4$ is disjoint from these paths and contains an $(\{v_3\}, X_5)$ -path, as desired. \square

Proposition 4.7.6. *If (B_0, \dots, B_5) is a $(t, 4)$ -chain in a graph G , and H is a connected subgraph of G such that each vertex in $V(H) \cap \bigcup_{i=0}^5 V(B_i)$ has degree 1 in H and $V(H) \cap V(B_3) \neq \emptyset$, then, for any $x_1 \in V(H) - \bigcup_{i=0}^5 V(B_i)$,*

$$\left(\bigcup_{i=1}^4 B_i \cup H, (V(B_0) \cap V(B_1)) \cup \{x_1\}, V(B_4) \cap V(B_5) \right)$$

is a t -bead; see Figure 4.13.

Proof. Let $\tilde{B} = \bigcup_{i=1}^4 B_i \cup H$. Let $X \subseteq (V(B_0) \cap V(B_1)) \cup \{x_1\}$ and $Y \subseteq V(B_4) \cap V(B_5)$ such that $|X| = |Y| \leq t$. Because \tilde{B} is connected, it suffices to show that \tilde{B} contains $|X|$ vertex-disjoint (X, Y) -paths. If $x_1 \notin X$ then $\bigcup_{i=1}^4 B_i$ contains $|X|$ vertex-disjoint (X, Y) paths, so we may assume $x_1 \in X$.




Figure 4.14: An illustration of Proposition 4.7.7. The vertices where H meets a bead only have one neighbor in H .

Let $v_3 \in V(H) \cap V(B_3)$. Note that H contains a path Q from x_1 to v_3 . The only vertices of Q that could have degree 1 in H are x_1 and v_3 . But $x_1 \notin \bigcup_{j=0}^5 V(B_j)$, so $V(Q) \cap \bigcup_{j=0}^5 V(B_j) = \{v_3\}$.

By Proposition 4.7.5, $\bigcup_{i=1}^4 B_i$ contains $|X|$ vertex-disjoint $((X - \{x_1\}) \cup \{v_1\}, Y)$ -paths; together with Q this gives $|X|$ vertex-disjoint (X, Y) -paths, as desired. \square

Proposition 4.7.7. *If (B_0, \dots, B_5) is a $(t, 4)$ -chain in a graph G , and H is a connected graph such that each vertex in $V(H) \cap \bigcup_{i=0}^5 V(B_i)$ has degree 1 in H and both $V(H) \cap V(B_2)$ and $V(H) \cap V(B_3)$ are non-empty, then, for any $x_1, x_4 \in V(H) - \bigcup_{i=0}^5 V(B_i)$,*

$$\left(\bigcup_{i=1}^4 B_i \cup H, (V(B_0) \cap V(B_1)) \cup \{x_1\}, (V(B_4) \cap V(B_5)) \cup \{x_4\} \right)$$

is a $(t + 1)$ -bead; see Figure 4.14.

Proof. Let $\tilde{B} = \bigcup_{i=1}^4 B_i \cup H$. Let $X \subseteq (V(B_0) \cap V(B_1)) \cup \{x_1\}$ and $Y \subseteq (V(B_4) \cap V(B_5)) \cup \{x_4\}$ such that $|X| = |Y| \leq t + 1$. Because \tilde{B} is connected, it suffices to show that \tilde{B} contains $|X|$ vertex-disjoint (X, Y) -paths.

If $x_4 \notin Y$ then $|Y| \leq |V(B_4) \cap V(B_5)| = t$, so by Proposition 4.7.6 $\bigcup_{i=1}^4 B_i \cup H$ contains $|X|$ vertex-disjoint (X, Y) -paths. Similarly, if $x_1 \notin X$ then $\bigcup_{i=1}^4 B_i \cup H$ contains $|X|$ vertex-disjoint (X, Y) -paths.

Otherwise $x_1 \in X$ and $x_4 \in Y$. Then $\bigcup_{i=1}^4 V(B_i)$ contains a collection \mathcal{P} of $|X| - 1$ vertex-disjoint $(X - \{x_1\}, Y - \{x_4\})$ -paths, and H contains a path Q from x_1 to x_4 . The only vertices in Q that could have degree 1 in H are x_1 and x_4 , but $x_1, x_4 \notin \bigcup_{i=0}^5 V(B_i)$.

Therefore, Q is an (x_1, x_4) -path disjoint from $\bigcup_{i=0}^5 V(B_i)$, and hence disjoint from each path in \mathcal{P} , as desired. \square

A jump in a necklace is closely related to the concept of a bridge: Note that a jump of a necklace is always contained in a bridge of the union of the beads in that necklace because the only two vertices of the jump that are in the necklace are the endpoints.

Lemma 4.7.8 shows that, if a necklace has a bridge with attachment vertices in each bead, then we can either find a hub or increase the connectivity of the necklace.

Lemma 4.7.8. *There exists a function $f_{4.7.8} : \mathbb{N} \rightarrow \mathbb{N}$ such that, for $t, s, m, n \in \mathbb{N}$ with $t \geq \max\{s, 1\}$, $n \geq 3$ and $m = f_{4.7.8}(n)$, if G is a graph containing a (t, s, m) -necklace $(B_i : i \in \mathbb{Z}_m)$, and H is a bridge of $\bigcup_{i=1}^m B_i$ in G such that, for each $i \in \{1, \dots, m\}$, H has an attachment vertex in $V(B_i)$, then either G contains a $(t+1, s, n)$ -necklace supported by $(B_i : i \in \mathbb{Z}_m)$ or G contains a (t, s, n) -necklace supported by $(B_i : i \in \mathbb{Z}_m)$ with a hub.*

Proof. Let $f_{4.7.8}(t, s, n) = \max\{5, f_{4.6.3}((4n+3)^2 + 3)\}$. Note that, because H is a bridge of $\bigcup_{i=1}^m V(B_i)$, $H - \bigcup_{i=1}^m V(B_i)$ is connected; choose a spanning tree T' of $H - \bigcup_{i=1}^m V(B_i)$. For each $i \in \{1, \dots, m\}$, let $v_i \in V(H) \cap V(B_i)$ and let $e_i \in E(H)$ be an edge from v_i to some vertex in $V(H) - \bigcup_{i=1}^m V(B_i)$; such an edge e_i exists because H has at least $\lceil m/2 \rceil \geq 3$ attachment vertices, so H is not a single edge, so H cannot contain any edge between two vertices of $\bigcup_{i=1}^m V(B_i)$. Let $T = T' \cup \{e_1, \dots, e_m\}$, so, for each $i \in \{1, \dots, m\}$, v_i is a leaf of T and $V(T) \cap \bigcup_{i=1}^m V(B_i) = \{v_1, \dots, v_m\}$.

Let $n' = (4n+3)^2 + 3$. Because $m \geq f_{4.6.3}(n')$, by Lemma 4.6.3, T contains a family $\{T_1, \dots, T_{n'}\}$ supported by $\{\{v_1\}, \dots, \{v_m\}\}$ such that either

- (a) T contains a hub z of $\{T_1, \dots, T_{n'}\}$ or,
- (b) for each $i \in \{1, \dots, n' - 1\}$, T contains an edge between T_i and T_{i+1} .

Choose such a family minimizing $|\bigcup_{i=1}^{n'} V(T_i)|$. If $i \in \{1, \dots, n'\}$ and $j, j' \in \{1, \dots, m\}$ such that $v_j, v_{j'} \in V(T_i)$ then v_j and $v_{j'}$ are both leaves of T , and hence leaves of $V(T_i)$, so removing $v_{j'}$ from $V(T_i)$ does not affect the external neighbours of T_i , so if $j \neq j'$, then T_i could be replaced by $T_i - v_{j'}$, contradicting minimality of $|\bigcup_{i=1}^{n'} V(T_i)|$. Therefore, for each $i \in \{1, \dots, n'\}$, $|V(T_i) \cap \{v_1, \dots, v_m\}| = 1$. For each $i \in \{1, \dots, n'\}$, let $a_i \in \{1, \dots, m\}$ such that $v_{a_i} \in V(T_i)$.

Suppose case (a) holds, so T contains a hub z of $\{T_1, \dots, T_{n'}\}$. By reordering $\{T_1, \dots, T_{n'}\}$, we may assume $a_1 < \dots < a_{n'}$. Note that $n \leq n'$. For each $i \in \{1, \dots, n-1\}$, let




Figure 4.15: Illustration of notation in the proof of Lemma 4.7.8

$B'_i = T_i \cup \bigcup_{j=a_i}^{a_{i+1}-1} B_j$ and let $B'_n = T_n \cup \bigcup_{j=a_n}^m \cup \bigcup_{j'=1}^{a_1-1} B_j$. By Lemma 4.5.6 and the fact that, for each i $T_i \cap \bigcup_{j=1}^m B_j = \{v_{a_i}\}$, we see that $(B'_i : i \in \mathbb{Z}_n)$ is a (t, s, n) -necklace supported by $(B_i : i \in \mathbb{Z}_m)$. z is a non-leaf in T , so $z \notin \bigcup_{i=1}^m V(B_i)$, so z is a hub of $(B'_i : i \in \mathbb{Z}_n)$, as desired.

Now suppose case (b) holds, so, for each $i \in \{1, \dots, n' - 1\}$, T contains an edge between T_i and T_{i+1} . Let $\tilde{n} = 4n + 4$. $n' - 2 = (\tilde{n} - 1)^2 + 1$, so by the Erdős-Szekeres Theorem, $(a_2, \dots, a_{n'-1})$ contains either an increasing subsequence of length \tilde{n} or a decreasing subsequence of length \tilde{n} . By reversing the order of $T_1, \dots, T_{n'}$ and $a_1, \dots, a_{n'}$ if necessary, we may assume $(a_2, \dots, a_{n'-1})$ contains an increasing subsequence $(\tilde{a}_1, \dots, \tilde{a}_{\tilde{n}})$.

For each $i \in \{1, \dots, \tilde{n}\}$, let $\tilde{T}_i = T_{\tilde{a}_i}$, and let \tilde{P}_i be a shortest path in T from $V(\tilde{T}_i)$ to $V(\tilde{T}_{i+1})$; note that $V(\tilde{P}_i) \subseteq \bigcup_{j=\tilde{a}_i}^{\tilde{a}_{i+1}-1} V(T_j)$. By minimality of \tilde{P}_i , $|V(\tilde{P}_i) \cap V(\tilde{T}_i)| = |V(\tilde{P}_i) \cap V(\tilde{T}_{i+1})| = 1$, so $\{\tilde{P}_1, \dots, \tilde{P}_{\tilde{n}}\}$ is a collection of internally vertex-disjoint paths. For each $i \in \{1, \dots, \tilde{n}\}$, $\tilde{P}_i \subseteq V(T)$ and does not contain any leaves of T , so $\tilde{P}_i \cap \bigcup_{j=1}^m V(B_j) = \emptyset$.

For each $i \in \{2, \dots, \tilde{n} - 1\}$, let $\tilde{B}_i = \bigcup_{j=\tilde{a}_i}^{\tilde{a}_{i+1}-1} B_j$ let $\tilde{B}_1 = \bigcup_{j=1}^{\tilde{a}_2-1} B_j$, and let $\tilde{B}_0 = \bigcup_{j=\tilde{a}_{\tilde{n}}}^m B_j$. For each $i \in \{1, \dots, \tilde{n} - 1\}$ let $\tilde{H}_i = \tilde{T}_i \cup \tilde{P}_i$ and let $\tilde{H}_{\tilde{n}} = \tilde{T}_{\tilde{n}}$. For each $i \in \{2, \dots, \tilde{n}\}$, let x_i be the unique vertex in $V(\tilde{P}_{i-1}) \cap V(\tilde{T}_i)$. For each $i \in \{1, \dots, n - 1\}$, let $B'_i = \bigcup_{j=4i+1}^{4i+4} (\tilde{B}_j \cup \tilde{H}_j)$ and let $B'_0 = \bigcup_{j=4n+1}^{4n+4} (\tilde{B}_j \cup \tilde{H}_j) \cup \bigcup_{j=1}^4 (\tilde{B}_j \cup \tilde{H}_j)$. See Figure 4.15

It now suffices to show that $(B'_i : i \in \mathbb{Z}_n)$ is a $(t + 1, s, n)$ -necklace supported by $(B_i : i \in \mathbb{Z}_m)$.

For each $i \in \mathbb{Z}_n$, $V(B'_i) \cap V(B'_{i-1}) = (V(B_{\tilde{a}_{4i}}) \cap V(B_{\tilde{a}_{4i+1}})) \cup \{x_{4i}\}$, which has size $t + 1$,

so necklace axiom (N1) holds.

For each $i, j \in \mathbb{Z}_n$, if $j - i \notin \{-1, 0, 1\}$, then $V(\tilde{H}_i) \subseteq \bigcup_{i'=\tilde{a}_i}^{\tilde{a}_{i+1}-1} V(T_{i'})$, which is disjoint from $\bigcup_{j'=\tilde{a}_j}^{\tilde{a}_{j+1}-1} V(T_{j'}) \supseteq V(\tilde{H}_j)$, so $V(B'_i) \cap V(B'_j) = \emptyset$, and axiom (N2) holds.

Note that, for each $i \in \{1, \dots, n-1\}$,

$$\left(\bigcup_{j=4i+1}^{4i+4} (\tilde{B}_j \cup \tilde{H}_j), V(\tilde{B}_{4i} \cup \tilde{H}_{4i}) \cap V(\tilde{B}_{4i+1} \cup \tilde{H}_{4i+1}), V(\tilde{B}_{4i+4} \cup \tilde{H}_{4i+4}) \cap V(\tilde{B}_{4i+5} \cup \tilde{H}_{4i+5}) \right)$$

is a $(t+1)$ -bead by Proposition 4.7.7, so axiom (N3) holds.

By Proposition 4.7.6,

$$\left(\bigcup_{j=4n+1}^{4n+4} (\tilde{B}_j \cup \tilde{H}_j), V(\tilde{B}_{4n} \cup \tilde{H}_{4n}) \cap V(\tilde{B}_{4n+1} \cup \tilde{H}_{4n+1}), V(\tilde{B}_{4n+4} \cup \tilde{H}_{4n+4}) \cap V(\tilde{B}_1 \cup \tilde{H}_1) \right)$$

is an s -bead; similarly,

$$\left(\bigcup_{j=1}^4 (\tilde{B}_j \cup \tilde{H}_j), V(\tilde{B}_{4n+4} \cup \tilde{H}_{4n+4}) \cap V(\tilde{B}_1 \cup \tilde{H}_1), V(\tilde{B}_4 \cup \tilde{H}_4) \cap V(\tilde{B}_5 \cup \tilde{H}_5) \right)$$

is an s -bead; thus,

$$(B'_0, V(B'_{n-1}) \cap V(B'_n), V(B'_n) \cap V(B'_1))$$

is an s -bead by Lemma 4.5.5. □

Lemma 4.7.9 shows that if a (t, s, n) -necklace has jumps to the weak bead from each other bead, then we can sacrifice the connectivity of the weak bead (decreasing s) in order to find a bridge with attachment vertices in each bead, which will allow us to recover the connectivity lost by decreasing s using Lemma 4.7.8.

Lemma 4.7.9. *There exists a function $f_{4.7.9} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that, for $t, s, m, n \in \mathbb{N}$ with $t \geq \max\{s, 1\}$, $n \geq 3$ and $m = f_{4.7.9}(s, n)$, if G is a graph containing a (t, s, m) -necklace $(B_i : i \in \mathbb{Z}_m)$ and, for each $i \in \{1, \dots, m-1\}$, G contains an (i, m) -jump of $(B_i : i \in \mathbb{Z}_m)$ then G contains either a $(t+1, \max\{s-1, 0\}, n)$ -necklace supported by $(B_i : i \in \mathbb{Z}_m)$ or a $(t, \max\{s-1, 0\}, n)$ -necklace supported by $(B_i : i \in \mathbb{Z}_m)$ and having a hub z .*

Proof. Let $f_{4.7.9}(s, n) = \max\{s, 2\}f_{4.7.8}(n) + 5$. For each $i \in \{2, \dots, m-2\}$, choose an (i, m) -jump J_i of $(B_i : i \in \mathbb{Z}_m)$ in G such that $|\bigcup_{i=2}^{m-2} E(J_i)|$ is minimized. Let $\tilde{m} = f_{4.7.8}(n)$.

By Lemma 4.7.8, it suffices to show that G has a $(t, \max\{s-1, 0\}, \tilde{m})$ -necklace $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{m}})$ supported by $(B_i : i \in \mathbb{Z}_m)$ and a bridge H of $\bigcup_{i=1}^{\tilde{m}} \tilde{B}_i$ such that, for each $i \in \{1, \dots, \tilde{m}\}$, H has an attachment vertex in \tilde{B}_i .

For $i \in \{2, \dots, m-2\}$, let v_i be the end of J_i in $V(B_i)$. Suppose $v \in V(\bigcup_{i=2}^{m-2} J_i) \cap V(\bigcup_{i=1}^m B_i)$ and $e_1, e_2 \in E(\bigcup_{j=2}^{m-2} J_j)$ are distinct edges incident with v . Let $I_1 = \{i \in \{2, \dots, m-2\} : e_1 \in E(J_i)\}$ and $I_2 = \{i \in \{2, \dots, m-2\} : e_2 \in E(J_i)\}$. For each $i_1 \in I_1$ and $i_2 \in I_2$, both J_{i_1} and J_{i_2} have one end at v and the other end in $V(B_m)$, so J_{i_1} is an (i_2, m) -jump. If, for each $i_2 \in I_2$, J_{i_2} is replaced by J_{i_1} for some $i_1 \in I_1$, then $|E(\bigcup_{j=2}^{m-2} J_j)|$ would be reduced, contradicting minimality. Therefore, for each $i \in \{2, \dots, m-2\}$, v_i has degree 1 in $\bigcup_{i=2}^{m-2} J_i$.

Suppose that $s = 0$. Let B_m^1 be the connected component of B_m containing $V(B_m) \cap V(B_1)$ and B_m^2 be the component of B_m containing $V(B_m) \cap V(B_{m-1})$; by (N6), $B_m = B_m^1 \cup B_m^2$. Let $I_1 = \{i \in \{2, \dots, m-2\} : V(J_i) \cap V(B_m^1) \neq \emptyset\}$ and let $I_2 = \{i \in \{2, \dots, m-2\} : V(J_i) \cap V(B_m^2) \neq \emptyset\}$. Note that $I_1 \cup I_2 = \{2, \dots, m-2\}$, so $\max\{|I_1|, |I_2|\} \geq (m-3)/2$; by possibly replacing $(B_i : i \in \mathbb{Z}_m)$ by $(B_{-i} : i \in \mathbb{Z}_m)$, we may assume $|I_1| \geq (m-3)/2 \geq \tilde{m}$. Let $H' = B_m^1 \cup \bigcup_{i \in I_1} J_i$, which is connected. Note that $V(H') \cap \bigcup_{i=2}^{m-2} V(B_i) = \{v_i : i \in I_1\}$ and for each $i \in I_1$, $v_i \notin B_m^1$, so v_i has degree 1 in H' . Hence, H' is contained in some bridge H of $\bigcup_{i=2}^{m-2} V(B_i)$.

Let $a_0, a_1, \dots, a_{\tilde{m}} \in I_1$ such that $2 \leq a_0 < a_1 < \dots < a_{\tilde{m}} \leq m-2$. For $i \in \{1, \dots, \tilde{m}-1\}$, let $\tilde{B}_i = \bigcup_{j=a_i}^{a_{i+1}-1} B_j$, and let $\tilde{B}_{\tilde{m}} = \bigcup_{j=a_{\tilde{m}}}^{m-2} B_j \cup \bigcup_{j=2}^{a_1-1} B_j$. Both $\bigcup_{j=a_{\tilde{m}}}^{m-2} B_j$ and $\bigcup_{j=2}^{a_1-1} B_j$ are connected, $V(\tilde{B}_{\tilde{m}}) \cap V(\tilde{B}_{\tilde{m}-1}) \subseteq V(B_{a_{\tilde{m}}})$ and $V(\tilde{B}_{\tilde{m}}) \cap V(\tilde{B}_1) \subseteq V(B_{a_1-1})$, so (N6) holds; (N1)-(N5) hold trivially, so $(\tilde{B}_1, \dots, \tilde{B}_{\tilde{m}})$ is a $(t, 0, \tilde{m})$ -necklace in G supported by (B_1, \dots, B_m) . H is a bridge of $\bigcup_{i=1}^{\tilde{m}} \tilde{B}_i$ and, for each $i \in \{1, \dots, \tilde{m}\}$, v_{a_i} is an attachment vertex of H in $V(\tilde{B}_i)$. Because $\tilde{m} = f_{4.7.8}(n)$, by Lemma 4.7.8, G contains either a $(t+1, 0, n)$ -necklace supported by $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{m}})$ or a $(t, 0, n)$ -necklace supported by $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{m}})$ and having a hub z , as desired.

Now suppose that $s > 0$. Choose a collection \mathcal{P} of s vertex-disjoint $(V(B_2) \cap V(B_1), V(B_{m-2}) \cap V(B_{m-1}))$ -paths in $B_1 \cup B_m \cup B_{m-1}$ minimizing $|\bigcup_{P \in \mathcal{P}} E(P)|$. By minimality of $|\bigcup_{P \in \mathcal{P}} E(P)|$, for each $P \in \mathcal{P}$, $V(P) \cap V(B_2)$ and $V(P) \cap V(B_{m-2})$ each contain exactly one vertex, which is an endpoint of P ; hence $P - (V(B_2) \cup V(B_{m-2}))$ is a path.

For each $i \in \{3, \dots, m-3\}$, let Q_i be a shortest path in B_m from $V(J_i) \cap V(B_m)$ to $\bigcup_{P \in \mathcal{P}} V(P) \cap V(B_m)$; by minimality, $|V(Q_i) \cap \bigcup_{P \in \mathcal{P}} V(P)| = 1$. For each $P \in \mathcal{P}$, let $I_P = \{i \in \{3, \dots, m-3\} : V(Q_i) \cap V(P) \neq \emptyset\}$. $\bigcup_{P \in \mathcal{P}} I_P = \{3, \dots, m-3\}$, so there is some $P_0 \in \mathcal{P}$ such that $|I_{P_0}| \geq (m-5)/s \geq \tilde{m}$. Let $a_0, a_1, \dots, a_{\tilde{m}} \in I_{P_0}$ such that $3 \leq a_1 < \dots < a_{\tilde{m}} \leq m-3$. For $i \in \{1, \dots, \tilde{m}-1\}$, let $\tilde{B}_i = \bigcup_{j=a_i}^{a_{i+1}-1} B_j$, and let $\tilde{B}_0 = \bigcup_{j=a_{\tilde{m}}}^{m-2} B_j \cup \bigcup_{P \in \mathcal{P} - \{P_0\}} \bigcup_{j=2}^{a_1-1} B_j$.

Claim 4.7.9.1. *We claim that $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{m}})$ a $(t, s - 1, \tilde{m})$ -necklace:*

Proof of Claim. For each $i \in \mathbb{Z}_{\tilde{m}}$, $V(\tilde{B}_i) \cap V(\tilde{B}_{i-1}) = V(B_{a_{i-1}}) \cap V(B_{a_i})$ which has size t , so necklace axiom (N1) holds.

Note that

$$\tilde{B}_0 \subseteq \bigcup_{j=a_{\tilde{m}}}^m B_j \cup \bigcup_{j=1}^{a_1-1} B_j,$$

so, for each $i \in \{1, \dots, \tilde{m}\}$, \tilde{B}_i is contained in the union of the beads between B_{a_i} and $B_{a_{i+1}-1}$; hence, if $j - i \notin \{-1, 0, 1\}$, then $\tilde{B}_i \cap \tilde{B}_j = \emptyset$, so axiom (N2) holds.

For each $i \in \{1, \dots, \tilde{m} - 1\}$, $\tilde{B}_i = \bigcup_{j=a_i}^{a_{i+1}-1} B_j$ and

$$\left(\bigcup_{j=a_i}^{a_{i+1}-1} B_j, V(B_{a_{i-1}}) \cap V(B_{a_i}), V(B_{a_{i+1}-1}) \cap V(B_{a_{i+1}}) \right)$$

is a t -bead by Lemma 4.5.5, so axiom (N3) holds.

To prove axiom (N4), note that \tilde{B}_0 is the union of two connected subgraphs, one containing $V(\tilde{B}_0) \cap V(\tilde{B}_1)$ and the other containing $V(\tilde{B}_0) \cap V(\tilde{B}_{\tilde{m}-1})$. Let $X \subseteq V(\tilde{B}_1) \cap V(\tilde{B}_{\tilde{m}}) = V(B_{a_1}) \cap V(B_{a_{\tilde{m}-1}})$ and $Y \subseteq V(\tilde{B}_{\tilde{m}}) \cap V(\tilde{B}_{\tilde{m}-1}) = V(B_{a_{\tilde{m}}}) \cap V(B_{a_{\tilde{m}-1}})$ such that $|X| = |Y| \leq s - 1$. It suffices to show that \tilde{B}_0 contains $|X|$ vertex-disjoint (X, Y) -paths.

Let $\mathcal{P}' \subseteq \mathcal{P} - \{P_0\}$ such that $|\mathcal{P}'| = |X|$. Let $X'_2 = V(B_2) \cap \bigcup_{P \in \mathcal{P}'} V(P)$ and $X'_{m-2} = V(B_{m-2}) \cap \bigcup_{P \in \mathcal{P}'} V(P)$, and note that $|X'_2| = |X'_{m-2}| = |X|$. Then $\bigcup_{j=2}^{\tilde{a}_1-1} B_j$ contains $|X|$ vertex-disjoint (X, X'_2) -paths, $\bigcup_{j=a_{\tilde{m}}}^{m-2} B_j$ contains $|X|$ vertex-disjoint (X'_{m-2}, Y) -paths, and \mathcal{P}' is a collection of $|X|$ vertex-disjoint (X'_2, X'_{m-2}) -paths that are internally vertex-disjoint from $\bigcup_{j=2}^{\tilde{a}_1-1} B_j \cup \bigcup_{j=\tilde{a}_{\tilde{m}}}^{m-2} B_j$. Therefore, \tilde{B}_0 contains $|X|$ vertex-disjoint (X, Y) -paths. \square (Claim)

Let $H' = (P_0 - (V(B_2) \cup V(B_{m-2}))) \cup \bigcup_{i \in I_{P_0}} (Q_i \cup J_i)$, which is connected. Note that H' is vertex-disjoint from $\bigcup_{P \in \mathcal{P} - \{P_0\}} P$, so $V(H') \cap \bigcup_{i=1}^{\tilde{m}} V(\tilde{B}_i) = V(H') \cap \bigcup_{i=3}^{m-3} V(B_i) = \{v_i : i \in I_{P_0}\}$. For each $i \in I_{P_0}$, $v_i \notin P_0 \cup \bigcup_{j \in I_{P_0}} Q_j$, and v_i has degree 1 in $\bigcup_{j=2}^{m-2} J_j$, so v_i has degree 1 in H' . Hence H' is contained in some bridge H of $\bigcup_{i=1}^{\tilde{m}} V(\tilde{B}_i)$, and, for each $i \in \{1, \dots, \tilde{m}\}$, v_i is an attachment vertex of H in \tilde{B}_i . Because $\tilde{m} \geq f_{4.7.8}(n)$, by Lemma 4.7.8, G contains either a $(t + 1, s - 1, n)$ -necklace supported by $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{m}})$ or a $(t, s - 1, n)$ -necklace supported by $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{m}})$ and having a hub z , as desired. \square

4.8 Growing a necklace

With Lemma 4.7.1, Lemma 4.7.4 and Lemma 4.7.9, we are now able to prove Lemma 4.8.1, showing that we can increase the connectivity of a necklace provided there is a surplus of paths between each pair of beads.

Lemma 4.8.1. *There exists a function $f_{4.8.1} : \mathbb{N}^3 \rightarrow \mathbb{N}$ such that, for $t, s, m, n \in \mathbb{N}$ with $t \geq \max\{s, 1\}$, $n \geq 3$ and $m = f_{4.8.1}(t, s, n)$, if G is a graph containing a (t, s, m) -necklace $(B_i : i \in \mathbb{Z}_m)$ and, for each $i, j \in \{1, \dots, m\}$, G contains $t + s + 1$ vertex-disjoint (B_i, B_j) -paths, then there exists $t', s', \ell' \in \mathbb{N}$ such that $t' + s' + \ell' = t + s + 1$ and G contains a (t', s', n) -necklace $(B'_i : i \in \mathbb{Z}_n)$ supported by $(B_i : i \in \mathbb{Z}_m)$ and a set Z' of ℓ' hubs of $(B'_i : i \in \mathbb{Z}_n)$.*

Proof. Let

$$\begin{aligned} n_1 &= f_{4.7.4}(t, s, n) \\ \hat{n} &= \begin{cases} n & \text{if } s = 0 \\ \max\{f_{4.8.1}(t + 1, s - 1, n), f_{4.8.1}(t, s - 1, n)\} & \text{if } s > 0 \end{cases} \\ n_2 &= f_{4.7.9}(s, \hat{n}) \\ \tilde{n} &= \max\{n_1, n_2\} \\ f_{4.8.1}(t, s, n) &= f_{4.7.1}(\tilde{n}) \end{aligned}$$

The proof goes by induction on s . By Lemma 4.7.1, either

1. G contains a (t, s, \tilde{n}) -necklace $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$ supported by $(B_i : i \in \mathbb{Z}_m)$ such that $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$ is jump-free in G or
2. G contains a (t, s, \tilde{n}) -necklace $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$ supported by $(B_i : i \in \mathbb{Z}_m)$ such that for each $i \in \{1, \dots, n - 1\}$, G contains an (i, \tilde{n}) -jump.

If case 1 holds then, by Lemma 4.7.4, there exists $t', s' \in \mathbb{N}$ such that $t' + s' = t + s + 1$ and G contains a (t', s', n) -necklace $(B'_i : i \in \mathbb{Z}_n)$ supported by $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$, and hence supported by $(B_i : i \in \mathbb{Z}_m)$, as desired.

Suppose case 2 holds. By Lemma 4.7.9, G contains a $(\hat{t}, \max\{s - 1, 0\}, \hat{n})$ -necklace $(\hat{B}_i : i \in \mathbb{Z}_{\hat{n}})$ supported by $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$ such that either $\hat{t} = t + 1$ or $\hat{t} = t$ and G contains a hub z of $(\hat{B}_i : i \in \mathbb{Z}_{\hat{n}})$. If $s = 0$ then $\hat{n} = n$ and $(B'_i : i \in \mathbb{Z}_n) = (\hat{B}_i : i \in \mathbb{Z}_{\hat{n}})$ is a $(\hat{t}, 0, n)$ -necklace such that either $\hat{t} = t + 1$ or z is a hub of $(B'_i : i \in \mathbb{Z}_n)$.

Suppose then that $s > 0$, so $(\hat{B}_i : i \in \mathbb{Z}_{\hat{n}})$ is a $(\hat{t}, s - 1, \hat{n})$ -necklace. If $\hat{t} = t + 1$ let $\hat{Z} = \emptyset$; if $\hat{t} = t$ let $\hat{Z} = \{z\}$, where z is a hub of $(\hat{B}_i : i \in \mathbb{Z}_{\hat{n}})$ in G ; let $\hat{\ell} = |\hat{Z}|$. For each $i, j \in \{1, \dots, \hat{n}\}$, G contains $t + s + 1$ vertex-disjoint (\hat{B}_i, \hat{B}_j) -paths, so $G - \hat{Z}$ contains $\hat{t} + s - 1$ vertex-disjoint (\hat{B}_i, \hat{B}_j) -paths. By the induction hypothesis, there exists $t', s', \ell' \in \mathbb{N}$ such that $t' + s' + \ell' = \hat{t} + (s - 1) + 1$ and $G - \hat{Z}$ contains a (t', s', n) -necklace $(B'_i : i \in \mathbb{Z}_n)$ and a set Z' of ℓ' hubs of $(B'_i : i \in \mathbb{Z}_n)$. Then $(B'_i : i \in \mathbb{Z}_n)$ is a (t', s', n) -necklace in G as well, and $Z' \cup \hat{Z}$ is a set of $\ell' + \hat{\ell}$ hubs of $(B'_i : i \in \mathbb{Z}_n)$ in G ; note that $t' + s' + \ell' + \hat{\ell} = \hat{t} + s + \hat{\ell} = t + s + 1$, as desired. \square

By an easy induction, it follows that, we can always increase the connectivity of a necklace up to the number of vertex-disjoint paths between the beads.

Corollary 4.8.2. *There exists a function $f_{4.8.2} : \mathbb{N}^4 \rightarrow \mathbb{N}$ such that, for $\theta, t, s, m, n \in \mathbb{N}$ with $\theta > 2$, $t \geq \max\{s, 1\}$, $n \geq 3$ and $m = f_{4.8.2}(\theta, t, s, n)$, if G is a graph containing a (t, s, m) -necklace $(B_i : i \in \mathbb{Z}_m)$ and, for each $i, j \in \{1, \dots, m\}$, G contains θ vertex-disjoint (B_i, B_j) -paths, then there exists $t', s', \ell' \in \mathbb{N}$ such that $t' + s' + \ell' = \theta$ and G contains a (t', s', n) -necklace $(B'_i : i \in \mathbb{Z}_n)$ supported by $(B_i : i \in \mathbb{Z}_m)$ and a set Z' of ℓ' hubs of $(B'_i : i \in \mathbb{Z}_n)$.*

Proof. By induction on $\theta - t - s$ using Lemma 4.8.1. \square

Now, finally, we are able to prove Lemma 4.5.2 (restated here for convenience).

Lemma 4.5.2. *There exists a function $f_{4.5.2} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that for each $\theta, m, n \in \mathbb{N}$ with $\theta \geq 3$, $n \geq 3$ and $m = f_{4.5.2}(\theta, n)$, if G is a graph containing a \mathcal{T} -tangle \mathcal{T} with covering number at least m , then either G contains a \mathcal{T} -aligned $K_{\theta, n}$ -minor or there exists $t, s, \ell \in \mathbb{N}$ with $t \geq s$, $t \geq 1$ and $t + s + \ell = \theta$ such that G contains a \mathcal{T} -aligned (t, s, n) -necklace with ℓ hubs.*

Proof. Let $n' = f_{4.8.2}(\theta, 1, 0, n)$ and let $f_{4.5.2}(\theta, n) = f_{4.6.5}(\theta, \max\{n, n'\})$. By Lemma 4.6.5, G contains either (a) a \mathcal{T} -aligned $K_{\theta, n}$ -minor, or (b) a \mathcal{T} -aligned $(1, 0, n')$ -necklace $(B'_i : i \in \mathbb{Z}_{n'})$ such that, for each $i, j \in \{1, \dots, n'\}$, G contains θ vertex-disjoint $(V(B'_i), V(B'_j))$ -paths. Case (a) is one of the desired outcomes, so we may assume that case (b) holds. By Corollary 4.8.2, there exists $t, s, \ell \in \mathbb{N}$ such that $t + s + \ell = \theta$ and G contains a (t, s, n) -necklace $(B_i : i \in \mathbb{Z}_n)$ supported by $(B'_i : i \in \mathbb{Z}_{n'})$ and a set Z of ℓ hubs of $(B_i : i \in \mathbb{Z}_n)$.

For each $i \in \{1, \dots, n'\}$, there is some $i' \in \{1, \dots, n'\}$ such that $E(B'_{i'}) \subseteq E(B_i)$; for each $A \in \mathcal{T}$, $E(B'_{i'}) \not\subseteq A$, so $E(B_i) \not\subseteq A$, so B_i is not 1-covered by \mathcal{T} . \square




Figure 4.16: An illustration of Lemma 4.9.1, showing the new weak bead B'_0 , containing the hub z , and showing how an additional path is routed through z in B'_0 .

4.9 Necklace to wheel

The final piece of the proof is to turn a necklace into a generalized-wheel, which we prove in this section. The biggest difference between a generalized-wheel and a necklace with hubs is that the necklace can have a weak bead, whereas the generalized-wheel has the same number of rim-matching edges all the way around; this is addressed in Lemma 4.9.1, Lemma 4.9.2 and Lemma 4.9.3

If a necklace has a hub and a weak bead, the hub can be sacrificed to increase the connectivity of the weak bead.

Lemma 4.9.1. *For each $\ell, t, s, n \in \mathbb{N}$ with $\ell > 0$, and $t > s$, if $m = n + 4$, and G is a graph containing a (t, s, m) -necklace $(B_i : i \in \mathbb{Z}_m)$ with a set Z of ℓ hubs then G contains a $(t, s + 1, n)$ -necklace $(B'_i : i \in \mathbb{Z}_n)$ with a set Z' of $\ell - 1$ hubs; moreover, for each $i \in \{1, \dots, n - 1\}$, $B'_i = B_{i+2}$; see Figure 4.16.*

Proof. Let $z \in Z$ and let e_2 be an edge from z to B_2 and let e_{m-2} be an edge from z to B_{m-2} . For each $i \in \{1, \dots, n-1\}$ let $B'_i = B_{i+2}$ and let $B'_0 = B_0 \cup B_1 \cup B_2 \cup B_{m-2} \cup B_{m-1} \cup \{e_2, e_{m-2}\}$. We claim that $(B'_i : i \in \mathbb{Z}_n)$ is a $(t, s + 1, n)$ -necklace. By Lemma 4.5.6 we know that

$$(B'_0 - \{e_2, e_{m-2}\}, B'_1, \dots, B'_{n-1})$$

is a (t, s, n) -necklace, and we also know that the only additional vertex in B'_0 is z , which

is not in any other bead. Therefore, it suffices to show that

$$(B'_0, V(B_2) \cap V(B_3), V(B_{m-2}) \cap V(B_{m-3}))$$

is an $(s + 1)$ -bead.

Let $X \subseteq V(B_2) \cap V(B_3)$ and $Y \subseteq V(B_{m-2}) \cap V(B_{m-3})$ such that $|X| = |Y| \leq s + 1$.

Choose a collection \mathcal{P}_2 of $|X|$ vertex-disjoint $(X, V(B_1) \cap V(B_m))$ -paths in $B_1 \cup B_2$ minimizing $|\bigcup_{P \in \mathcal{P}_2} E(P)|$; by minimality, for each $P \in \mathcal{P}_2$, P intersects $V(B_0)$ only at an endpoint. Let $v_2 \in V(B_2)$ be the vertex incident with e_2 . Choose a shortest path Q_2 in B_2 from v_2 to $\bigcup_{P \in \mathcal{P}_2} V(P)$; by minimality, $|V(Q_2) \cap \bigcup_{P \in \mathcal{P}_2} V(P)| = 1$. Let $P_2 \in \mathcal{P}_2$ be the path for which $V(P_2) \cap V(Q_2) \neq \emptyset$. Let P'_2 be the subpath of P_2 from the end in X to v_2 , together with Q_2 and e_2 , so P'_2 is an $(X, \{z\})$ -path; moreover, $v_2 \in V(B_2)$ so $v_2 \notin V(B_0)$, so $P'_2 \cap V(B_m) = \emptyset$. Let $X_1 = V(B_0) \cap \bigcup_{P \in \mathcal{P}_2 - \{P_2\}} V(P)$ and note that $|X_1| = |X| - 1$.

Similarly, choose a collection \mathcal{P}_{m-2} of $|X|$ vertex-disjoint $(Y, V(B_{m-1}) \cap V(B_0))$ -paths in $B_{m-1} \cup B_{m-2}$ minimizing $|\bigcup_{P \in \mathcal{P}_{m-2}} E(P)|$. Let $v_{m-2} \in V(B_{m-2})$ be the vertex incident with e_{m-2} . Choose a shortest path Q_{m-2} in B_{m-2} from v_{m-2} to $\bigcup_{P \in \mathcal{P}_{m-2}} V(P)$. Let $P_{m-2} \in \mathcal{P}_{m-2}$ be the path for which $V(P_{m-2}) \cap V(Q_{m-2}) \neq \emptyset$. Let P'_{m-2} be the subpath of P_{m-2} from the end in Y to v_{m-2} , together with Q_{m-2} and e_{m-2} , so P'_{m-2} . Let $X_{m-1} = V(B_0) \cap \bigcup_{P \in \mathcal{P}_{m-2} - \{P_{m-2}\}} V(P)$.

Let \mathcal{P}_0 be a collection of $|X_1| = |X| - 1 \leq s$ vertex-disjoint (X_1, X_{m-1}) -paths in B_0 . Then $\bigcup(\mathcal{P}_2 - \{P_2\}) \cup \bigcup \mathcal{P}_0 \cup (\mathcal{P}_{m-2} - \{P_{m-2}\})$ is the union of $|X| - 1$ vertex-disjoint (X, Y) -paths. $P'_2 \cup P'_{m-2}$ is another (X, Y) -path disjoint from all of these, giving $|X|$ vertex-disjoint (X, Y) -paths. Therefore,

$$(B'_0, V(B'_0) \cap V(B'_1), V(B'_0) \cap V(B'_{n-1}))$$

is an $(s + 1)$ -bead, so $(B'_i : i \in \mathbb{Z}_n)$ is a $(t, s + 1, n)$ -necklace in G , and $Z' = Z - \{z\}$ is a set of $\ell - 1$ hubs of $(B'_i : i \in \mathbb{Z}_n)$, as desired. \square

If we run out of hubs to sacrifice with Lemma 4.9.1, but still have a weak bead, we can sacrifice the connectivity of all of the other beads to produce one more hub.

Lemma 4.9.2. *There exists a function $f_{4.9.2} : \mathbb{N} \rightarrow \mathbb{N}$ such that, for $t, s, m, n \in \mathbb{N}$ with $t > \max\{s, 1\}$, $n \geq 3$ and $m = f_{4.9.2}(t, n)$, if G is a graph containing a (t, s, m) -necklace $(B_i : i \in \mathbb{Z}_m)$ then G contains a minor G' which contains a $(t - 1, s, n)$ -necklace $(B'_i : i \in \mathbb{Z}_n)$ with a hub z' ; moreover, there exists $c_1, \dots, c_n \in \{1, \dots, m\}$ such that $c_1 < \dots < c_n$ and, for each $i \in \{1, \dots, n - 1\}$, $E(B'_i) \subseteq \bigcup_{j=c_i}^{c_{i+1}-1} E(B_j)$; see Figure 4.17.*

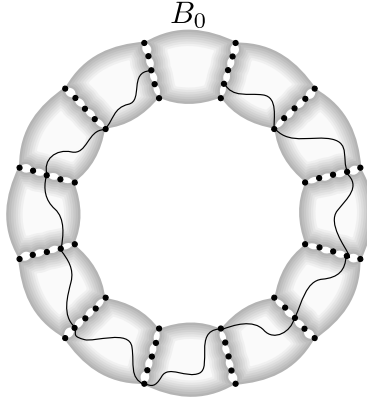


Figure 4.17: An illustration of Lemma 4.9.2. A path through all of the non-weak-beads can be contracted to a single vertex, forming a hub. This path must be chosen carefully to avoid breaking too many of the beads in the necklace.

Proof. Let $f_{4.9.2}(t, n) = 2t^2(n - 1) + 1$. By Lemma 4.5.5,

$$\left(\bigcup_{i=1}^{m-1} B_i, V(B_1) \cap V(B_0), V(B_{m-1}) \cap V(B_0) \right)$$

is a t -bead, so $\bigcup_{i=1}^{m-1} B_i$ contains a family \mathcal{P} of t vertex-disjoint $(V(B_1) \cap V(B_0), V(B_{m-1}) \cap V(B_0))$ -paths. Note that, for each $i \in \{1, \dots, m - 1\}$, there exists $P_i \in \mathcal{P}$ such that the paths in $\mathcal{P} - \{P_i\}$ all lie in a single connected component of $B_i - V(P_i)$. Therefore, there exists some $P_0 \in \mathcal{P}$ such that

$$|\{i \in \{1, \dots, m - 1\} : P_i = P_0\}| \geq \frac{m - 1}{t} = 2t(n - 1).$$

Let

$$I_0 = \{i \in \{1, \dots, m - 1\} : i \equiv 0 \pmod{2} \wedge P_i = P_0\}$$

and

$$I_1 = \{i \in \{1, \dots, m - 1\} : i \equiv 1 \pmod{2} \wedge P_i = P_0\},$$

so either $|I_0| \geq t(n - 1)$ or $|I_1| \geq t(n - 1)$. In either case there exists a sequence $(a_1, \dots, a_{t(n-1)})$ of integers in $\{1, \dots, m - 1\}$ such that for each $i \in \{1, \dots, t(n - 1) - 1\}$, $a_i + 1 < a_{i+1}$, and, for each $i \in \{1, \dots, t(n - 1)\}$, $P_i = P_0$. For $i \in \{1, \dots, t(n - 1)\}$, let \tilde{B}_i be the connected component of $B_{a_i} - V(P_0)$ containing all of the paths in $\mathcal{P} - \{P_0\}$; let $\tilde{B}_0 = V(B_1) \cap V(B_0)$.

For $i \in \{1, \dots, n-1\}$, let \mathcal{P}'_i be the set containing, for each $P \in \mathcal{P} - \{P_0\}$, a shortest subpath of P from $V(\tilde{B}_{t(i-1)})$ to $V(\tilde{B}_{ti})$ and let $B'_i = \bigcup_{j=t(i-1)+1}^{ti} \tilde{B}_j \cup \bigcup \mathcal{P}'_i$; let $B'_0 = \bigcup_{j=(a_{t(n-1)}+1)}^m B_j$. For $i \in \{1, \dots, n\}$, let $c_i = a_{t(i-1)+1}$ and note that, for $i \in \{1, \dots, n-1\}$, $E(B'_i) \subseteq \bigcup_{j=c_i}^{c_{i+1}-1} E(B_j)$. We claim that $(B'_i : i \in \mathbb{Z}_n)$ is a $(t-1, s, n)$ -necklace. The only non-trivial facts to verify are (N1), that, for each $i \in \{1, \dots, n\}$, $|V(B'_i) \cap V(B'_{i+1})| = t-1$ and (N3), that, for each $i \in \{1, \dots, n-1\}$, $(B'_i, V(B'_{i-1}) \cap V(B'_i), V(B'_i) \cap V(B'_{i+1}))$ is a $(t-1)$ -bead.

For each $i \in \{1, \dots, n-1\}$, because the paths in \mathcal{P}'_i are shortest paths, they each intersect $V(\tilde{B}_{t(i-1)})$ in exactly one vertex, and that vertex is in $(V(B_{a_{t(i-1)}}) \cap V(B_{a_{t(i-1)}+1})) - V(P_0)$, which has size $t-1 = |\mathcal{P}'_i|$; hence $|V(B'_i) \cap V(B'_{i-1})| = t-1$. Also $V(B'_0) \cap V(B'_{n-1}) = (V(B_{a_{t(n-1)}+1}) \cap V(B_{a_{t(n-1)}})) - V(P_0)$, which has size $t-1$. Therefore (N1) holds.

To prove (N3), suppose $i \in \{1, \dots, n-1\}$, $X \subseteq V(B'_{i-1}) \cap V(B'_i) = V(B_{a_{t(i-1)}}) \cap \bigcup_{P \in \mathcal{P}'_i} V(P)$ and $Y \subseteq V(B'_i) \cap V(B'_{i+1}) = V(B_{a_{ti+1}}) \cap \bigcup_{P \in \mathcal{P}'_i} V(P)$ such that $|X| = |Y| \leq t-1$. By Lemma 4.7.2 with $\mathcal{A} = \{P \in \mathcal{P}'_i : V(P) \cap X \neq \emptyset\}$, $\mathcal{B} = \{P \in \mathcal{P}'_i : V(P) \cap Y \neq \emptyset\}$ and $\mathcal{H} = \{\tilde{B}_1, \dots, \tilde{B}_{t(n-1)}\}$, B'_i contains $|X|$ vertex-disjoint (X, Y) -paths, as desired. Hence, $(B'_i : i \in \mathbb{Z}_n)$ is a $(t-1, s, n)$ -necklace in G .

Let P'_0 be the shortest $(V(B_{a_1}), V(B_{a_{t(n-1)}+1}))$ -subpath of P_0 ; by minimality, P'_0 is internally vertex-disjoint from $V(B_{a_1})$ and $V(B_{a_{t(n-1)}+1})$. Therefore, P'_0 is internally vertex-disjoint from $\bigcup_{i=1}^n V(B'_i)$, and its endpoint in $V(B_{a_{t(n-1)}+1})$ is in B'_0 . Let e' be the edge of P'_0 incident with the endpoint in $V(B'_0)$.

Note that, for each $i \in \{1, \dots, n\}$, there is an edge from $P'_0 - \{e'\}$ to B'_i . Thus, if G' is obtained from G by contracting $P'_0 - \{e'\}$ down to a single vertex z' , then $(B'_i : i \in \mathbb{Z}_n)$ is a $(t-1, s, n)$ -necklace in G' with a hub, z' , as desired. \square

By repeatedly applying Lemma 4.9.1 or Lemma 4.9.2, we can arrive at a θ -linked necklace with no weak bead.

Lemma 4.9.3. *There exists a function $f_{4.9.3} : \mathbb{N}^4 \rightarrow \mathbb{N}$ such that, for $\ell, t, s, m, n \in \mathbb{N}$ with $t \geq \max\{s, 1\}$, $t + s + \ell > 1$, $n \geq 3$ and $m = f_{4.9.3}(\ell, t, s, n)$, if G is a graph containing a (t, s, m) -necklace $(B_i : i \in \mathbb{Z}_m)$ with a set Z of ℓ hubs then there exists $t', \ell' \in \mathbb{N}$ such that $2t' + \ell' = t + s + \ell$ and G contains a minor G' which contains a (t', t', n) -necklace $(B'_i : i \in \mathbb{Z}_n)$ with a set Z' of ℓ' hubs; moreover, there exists $c_1, \dots, c_n \in \{1, \dots, m\}$ such that $c_1 < \dots < c_n$ and, for each $i \in \{1, \dots, n-1\}$, $E(B'_i) \subseteq \bigcup_{j=c_i}^{c_{i+1}-1} E(B_j)$.*

Proof. By induction on $t - s$. If $t = s$ then let $f_{4.9.3}(\ell, t, t, n) = n$, so $(B_i : i \in \mathbb{Z}_m)$ is a (t, t, n) -necklace, so the lemma holds in this case.

Suppose $t > s$ and $\ell > 0$. In this case let $\tilde{n} = f_{4.9.3}(\ell - 1, t, s + 1, n)$ and let $f_{4.9.3}(\ell, t, s, n) = \tilde{n} + 4$. By Lemma 4.9.1, G contains a $(t, s + 1, \tilde{n})$ -necklace $(\tilde{B}_i : i \in \mathbb{Z}_{\tilde{n}})$ and a set \tilde{Z} of $\ell - 1$ hubs. By the induction hypothesis, there exists a minor G' of G containing a (t', t', n) -necklace $(B'_i : i \in \mathbb{Z}_n)$ with a set Z' of ℓ' hubs, where $2t' + \ell' = t + (s + 1) + (\ell - 1) = t + s + \ell$. Moreover, there exists $\tilde{c}_1, \dots, \tilde{c}_n \in \{1, \dots, \tilde{n}\}$ such that $\tilde{c}_1 < \dots < \tilde{c}_n$ and, for each $i \in \{1, \dots, n\}$, $\bigcup E(B'_i) \subseteq \bigcup_{j=\tilde{c}_i}^{\tilde{c}_{i+1}-1} E(\tilde{B}_j)$. But, by Lemma 4.9.1, for each $i \in \{1, \dots, \tilde{n} - 1\}$, $\tilde{B}_i = B_{i+2}$, so, if $c_i = \tilde{c}_i + 2$ then, for each $i \in \{1, \dots, n - 1\}$, $E(B'_i) \subseteq \bigcup_{j=c_i}^{c_{i+1}-1} E(B_j)$, as desired. Thus, the lemma holds in this case.

Finally suppose $t > s$ and $\ell = 0$. Note that $1 < t + s + \ell \leq 2t - 1$ so $t > 1$. In this case let $\hat{n} = f_{4.9.3}(1, t - 1, s, n)$ and let $f_{4.9.3}(\ell, t, s, n) = f_{4.9.2}(t, \hat{n})$. By Lemma 4.9.2, G contains a minor \hat{G} which contains a $(t - 1, s, \hat{n})$ -necklace $(\hat{B}_i : i \in \mathbb{Z}_{\hat{n}})$ and a hub \hat{z} . By the induction hypothesis, there exists a minor G' of \hat{G} containing a (t', t', n) -necklace $(B'_i : i \in \mathbb{Z}_n)$ with a set Z' of ℓ' hubs, where $2t' + \ell' = (t - 1) + s + 1 = t + s + \ell$. Moreover, there exists $c'_1, \dots, c'_n \in \{1, \dots, \hat{n}\}$ such that $c'_1 < \dots < c'_n$ and, for each $i \in \{1, \dots, n - 1\}$, $E(B'_i) \subseteq \bigcup_{j=c'_i}^{c'_{i+1}-1} E(\hat{B}_j)$. By Lemma 4.9.2, there exists $\hat{c}_1, \dots, \hat{c}_{\hat{n}}$ such that $\hat{c}_1 < \dots < \hat{c}_{\hat{n}}$ and for each $i \in \{1, \dots, \hat{n} - 1\}$, $E(\hat{B}_i) \subseteq \bigcup_{j=\hat{c}_i}^{\hat{c}_{i+1}-1} E(B_j)$. Then, for each $i \in \{1, \dots, n - 1\}$, $E(B'_i) \subseteq \bigcup_{j=c'_i}^{c'_{i+1}-1} \bigcup_{k=\hat{c}_j}^{\hat{c}_{j+1}-1} E(B_k)$. For $i \in \{1, \dots, n\}$, let $c_i = \hat{c}_{c'_i}$. Therefore, for $i \in \{1, \dots, n - 1\}$, $E(B'_i) \subseteq \bigcup_{j=c_i}^{c_{i+1}-1} E(B_j)$, as desired. Thus, the lemma holds. \square

Once we have a necklace with no weak bead, obtaining a wheel-minor is straightforward.

Lemma 4.9.4. *There exists a function $f_{4.9.4} : \mathbb{N}^4 \rightarrow \mathbb{N}$ such that, for $\ell, t, s, m, n \in \mathbb{N}$ with $t \geq \max\{s, 1\}$, $t + s + \ell > 1$, $n > 2(t + s + \ell - 1)$ and $m = f_{4.9.4}(\ell, t, s, n)$, if G is a graph containing a $(t + s + \ell)$ -tangle \mathcal{T} , a \mathcal{T} -aligned (t, s, m) -necklace $(B_i : i \in \mathbb{Z}_m)$, and a set Z of ℓ hubs of $(B_i : i \in \mathbb{Z}_m)$, then there exists $r, \ell' \in \mathbb{N}$ such that $2r + \ell' = \ell + t + s$ and G contains an \mathcal{T} -aligned (r, ℓ', n) -wheel-minor.*

Proof. Let $m = f_{4.9.4}(\ell, t, s, n) = f_{4.9.3}(\ell, t, s, 2n)$. By Lemma 4.9.3, G contains a minor \tilde{G} which contains a $(\tilde{t}, \tilde{t}, 2n)$ -necklace $(\tilde{B}_i : i \in \mathbb{Z}_{2n})$ with a set \tilde{Z} of $\tilde{\ell}$ hubs and $2\tilde{t} + \tilde{\ell} = t + s + \ell$. Moreover, there exists $c_1, \dots, c_{2n} \in \{1, \dots, m\}$ such that $c_1 < \dots < c_{2n}$ and, for each $i \in \{1, \dots, 2n - 1\}$, $E(\tilde{B}_i) \subseteq \bigcup_{j=c_i}^{c_{i+1}-1} E(B_j)$.

For each $i \in \{1, \dots, 2n\}$, let \mathcal{P}_i be a collection of \tilde{t} vertex-disjoint paths between $V(\tilde{B}_{i-1}) \cap V(\tilde{B}_i)$ and $V(\tilde{B}_i) \cap V(\tilde{B}_{i+1})$; because $(\tilde{B}_i : i \in \mathbb{Z}_{2n})$ is a $(\tilde{t}, \tilde{t}, 2n)$ -necklace,

$$|V(\tilde{B}_{i-1}) \cap V(\tilde{B}_i)| = |V(\tilde{B}_i) \cap V(\tilde{B}_{i+1})| = \tilde{t} = |\mathcal{P}_i|,$$

so, for each $P \in \mathcal{P}_i$,

$$|V(P) \cap V(\tilde{B}_{i-1})| = |V(P) \cap V(\tilde{B}_{i+1})| = 1.$$

Notice that, because $V(B_{i-1}) \cap V(B_{i+1}) = \emptyset$, each $P \in \mathcal{P}_i$ has at least one edge, $e_{P,i}$.

For each $i \in \{1, \dots, 2n\}$, \tilde{B}_i is connected, so $V(\tilde{B}_i)$ can be partitioned into $\{X_{i,P} : P \in \mathcal{P}_i\}$ such that, for each $P \in \mathcal{P}_i$, $V(P) \subseteq X_{i,P}$ and $X_{i,P}$ is connected; let $T_{i,P}$ be a spanning tree of $X_{i,P}$. Let

$$G' = \tilde{G} / \bigcup_{i=1}^n \left(\bigcup_{P \in \mathcal{P}_{2i-1}} E(T_{2i-1,P}) \cup \bigcup_{P \in \mathcal{P}_{2i}} E(P) - \{e_{P,i}\} \right).$$

For each $i \in \mathbb{Z}_n$, \tilde{B}_{2i-1} gets contracted to a connected subgraph of G' with exactly $|\mathcal{P}_i| = t$ vertices, which has a spanning tree T_i . For each $i \in \mathbb{Z}_n$, \tilde{B}_{2i} gets contracted to a subgraph of G' containing a perfect matching, M_i , between T_i and T_{i+1} . For each $i \in \mathbb{Z}_n$ and each $z \in \tilde{Z}$, the neighbour of z in \tilde{B}_{2i-1} in \tilde{G} is contracted onto a vertex of T_i , so z is a hub of $\{T_1, \dots, T_i\}$ in G' ; for $i \in \mathbb{Z}_n$ and $z \in Z$, let $e_{i,z}$ be an edge incident with z and a vertex in T_i . Thus,

$$\mathcal{W} = ((T_i : i \in \mathbb{Z}_n), (M_i : i \in \mathbb{Z}_n), Z, (e_{i,z} : i \in \mathbb{Z}_n, z \in Z))$$

is a model for a $(\tilde{t}, \tilde{\ell}, n)$ -wheel, so G' has a $(\tilde{t}, \tilde{\ell}, n)$ -wheel subgraph, W , and G contains a $(\tilde{t}, \tilde{\ell}, n)$ -wheel-minor.

We must also show that W is \mathcal{T} -aligned; that is, we must show that the fundamental tangle of \mathcal{W} , \mathcal{T}' , induces a subset of the tangle \mathcal{T} in G . Let $\theta = t + s + \ell$. Let (X, Y) be a $(\theta - 1)$ -separation in G , so $(X \cap E(W), Y \cap E(W))$ is a $(\theta - 1)$ -separation in W .

Suppose that $X \cap E(W) \in \mathcal{T}'$. Let

$$I = \{i \in \mathbb{Z}_n : X \cap E(T_i^+) \neq \emptyset\}.$$

By Lemma 4.3.2, $|I| \leq \theta - 1 < n/2$.

Claim 4.9.4.1. $X \in \mathcal{T}$.

Proof of Claim. Let $J = \{1, \dots, n\} - I$ and note that, for each $j \in J$, $E(T'_j) \subseteq Y$. Because $|I| < \theta$, $|J| > n - \theta > \theta$. For $j \in J$, the rim-matching between T_j and T_{j+1} is a collection of $\tilde{t} > 1$ edges in \tilde{B}_{2j} , each of which are in Y . In particular, Y contains

an edge in \tilde{B}_{2j} , which is an edge in $E\left(\bigcup_{i=c_{2j}}^{c_{2j+1}-1} B_i\right)$. Hence, for $j \in J$, if $\bigcup_{i=c_{2j}}^{c_{2j+1}-1} B_i$ contains an edge in X then, because $\bigcup_{i=c_{2j}}^{c_{2j+1}-1} B_i$ is connected, this subgraph must contain a vertex in $\partial(X)$. For $j, j' \in \{1, \dots, n\}$, if $j < j'$ then $c_{2j+1} < c_{2j'}$, so $\bigcup_{i=c_{2j}}^{c_{2j+1}-1} B_i$ and $\bigcup_{i=c_{2j'}}^{c_{2j'+1}-1} B_i$ are vertex-disjoint. Therefore, because $|J| > \theta$, there is some $j_0 \in J$ such that $E\left(\bigcup_{i=c_{2j_0}}^{c_{2j_0+1}-1} B_i\right) \cap X = \emptyset$. Therefore, $E(B_{c_{2j_0}}) \subseteq Y$, so, because $B_{c_{2j_0}}$ is not 1-covered by \mathcal{T} , $Y \notin \mathcal{T}$. Hence, $X \in \mathcal{T}$, proving the claim. \square (Claim)

Thus, \mathcal{T} is the tangle induced in G by $\mathcal{T}_\theta(W)$, so W is \mathcal{T} -aligned, proving the lemma. \square

Now that we can go from a tangle to a necklace, to a generalized-wheel-minor, the proof of Theorem 4.4.1 (restated here) is immediate.

Theorem 4.4.1. *There exists a function $f_{4.4.1} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that for $\theta, n, N \in \mathbb{N}$ with $\theta \geq 2$, $n > 2(\theta - 1)$, and $N = f_{4.4.1}(\theta, n)$, if G is a graph containing a θ -tangle \mathcal{T} with covering-number at least N , then either G contains a \mathcal{T} -aligned $K_{\theta, n}$ -minor or there exists $r, \ell \in \mathbb{N}$ such that $2r + \ell = \theta$ and G contains a \mathcal{T} -aligned (r, ℓ, n) -wheel-minor.*

Proof. Let

$$\tilde{n} = \max_{\ell+t+s=\theta} f_{4.9.4}(\ell, t, s, n)$$

and let $f_{4.4.1}(\theta, n) = f_{4.5.2}(\theta, \tilde{n})$. By Lemma 4.5.2, there exists $t, s, \ell \in \mathbb{N}$ such that $t+s+\ell = \theta$ and G contains a (t, s, \tilde{n}) -necklace $(B_i : i \in \mathbb{Z}_{\tilde{n}})$ and a set Z of hubs of $(B_i : i \in \mathbb{Z}_{\tilde{n}})$. By Lemma 4.9.4, there exists $r, \ell' \in \mathbb{N}$ such that $2r + \ell' = t + s + \ell = \theta$ and G contains a (r, ℓ', n) -wheel-minor. \square

4.10 Grid Theorem

To conclude this chapter, we will prove Theorem 4.4.2 and use this to prove the Grid Theorem (Theorem 4.4.3).

4.10.1 Constructing homogeneous wheels

This section proves Theorem 4.4.2, which is a slight refinement of Theorem 4.4.1 that is useful in some applications, including our proof of the grid theorem.

Theorem 4.4.2. *There exists a function $f_{4.4.2} : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that, for each $\theta, n, N \in \mathbb{N}$, with $\theta > 1$, $n > 2(\theta - 1)$ and $N = f_{4.4.2}(\theta, n)$, if G is a graph containing a θ -tangle \mathcal{T} with covering-number at least N , then either G contains a \mathcal{T} -aligned $K_{\theta, n}$ -minor or there exists $r, \ell \in \mathbb{N}$ such that $2r + \ell = \theta$ and G contains a \mathcal{T} -aligned, homogeneous (r, ℓ, n) -wheel-minor.*

Proof. Let $m' = \max\{nr^{r+\ell-2} : r, \ell \in \mathbb{N}, 2r + \ell = \theta\}$ and let $f_{4.4.2}(\theta, n) = f_{4.4.1}(\theta, m')$. By Theorem 4.4.1, G contains an (r, ℓ, n) -wheel-minor G' with model

$$\mathcal{W} = ((T_i : i \in \mathbb{Z}_{m'}), (M_i : i \in \mathbb{Z}_{m'}), Z, (e_{i,z} : i \in \mathbb{Z}_{m'}, z \in Z)).$$

Note that $\bigcup_{i=1}^{m'-1} M_i$ is the union of r vertex-disjoint paths, P_1, \dots, P_r . For $i \in \mathbb{Z}_{m'}$ and $j \in \{1, \dots, r\}$, $|V(T_i) \cap V(P_j)| = 1$, so there is a tree T'_i with $V(T'_i) = \{1, \dots, r\}$ and an isomorphism ψ_i from T_i to T'_i where $\psi_i(v) = j$ if $V(T_i) \cap V(P_j) = \{v\}$. For $i \in \mathbb{Z}'_m$ and $z \in Z$, let $a_{i,z} \in \{1, \dots, r\}$ such that $\psi_i^{-1}(a_{i,z})$ is incident with $e_{i,z}$.

By Cayley's Formula, there are r^{r-2} trees with vertex set $\{1, \dots, r\}$. There are r^ℓ ways to choose a sequence of ℓ values in $\{1, \dots, r\}$. Therefore, because $m' \geq nr^{r+\ell-2}$, there is some tree T' with $V(T') = \{1, \dots, r\}$, some $b_z \in \{1, \dots, r\}$ for each $z \in Z$, and some $I \subseteq \{1, \dots, m'\}$ such that for each $i \in I$, $T'_i = T'$ and $a_{i,z} = b_z$ for each $z \in Z$, and $|I| = n$.

Let \hat{G} be obtained from G' by contracting each edge in $\bigcup_{i \in \{1, \dots, m'\} - I} M_i$. Let i_1, \dots, i_n be the elements of I where $i_1 < \dots < i_n$. Then \hat{G} is isomorphic to a homogeneous (r, ℓ, n) -wheel with model $\hat{\mathcal{W}}$; moreover, \hat{G} is \mathcal{T} -aligned in G because the fundamental tangle of $\hat{\mathcal{W}}$ and the fundamental tangle of \mathcal{W} both induce the same θ -tangle in G , and that tangle must be contained in \mathcal{T} by Theorem 4.4.1. \square

4.10.2 Tangle truncation

In this subsection we present a general tool for tangles that has not yet come up in this thesis but is needed for our proof of the Grid Theorem. If \mathcal{T} is a θ -tangle in a connectivity system (S, λ) , and $\theta' \leq \theta$, then $\mathcal{T}' = \{A \in \mathcal{T} : \lambda(A) < \theta'\}$ is the *truncation of \mathcal{T} to order θ'* . If each $x \in S$ is contained in some $A_x \in \mathcal{T}$ with $\lambda(A_x) < \theta'$, then \mathcal{T}' is a tangle. In the case of graphs, for each $e \in E(G)$, $\lambda_G(\{e\}) < 3$, so if \mathcal{T} is a θ -tangle in the graph G and $3 \leq \theta' \leq \theta$, then the truncation of \mathcal{T} to order θ' is a tangle in G .

If a tangle is truncated to a much smaller order, the covering number must increase.

Lemma 4.10.1. *For $\theta, \theta' \in \mathbb{N}$ with $2\theta' \leq \theta$, if \mathcal{T} is a θ -tangle with covering-number δ , \mathcal{T}' is the truncation of \mathcal{T} to order θ' , and δ' is the covering-number of \mathcal{T}' , then $\delta' \geq \delta + 1$.*

Proof. Suppose $\{A_1, \dots, A_\delta\} \subseteq \mathcal{T}$ and $\bigcup_{i=1}^\delta A_i = E(G)$. But $A_1 \cup A_2 \notin \mathcal{T}$ or else $\{A_1 \cup A_2, A_3, \dots, A_\delta\}$ would be a cover of $E(G)$ by $\delta - 1$ sets in \mathcal{T} . Therefore, $2\theta' \leq \theta \leq \lambda(A_1 \cup A_2) \leq \lambda(A_1) + \lambda(A_2)$, so either $\lambda(A_1) \geq \theta'$ or $\lambda(A_2) \geq \theta'$. Therefore, $E(G)$ cannot be covered by δ sets in \mathcal{T} each having connectivity less than θ' , so $\delta' > \delta$, proving the lemma. \square

4.10.3 Proof of Grid Theorem

Lemma 4.10.2. *For $a, b, n \in \mathbb{N}$ with $n \geq 3$, $a, b \geq n^2$, if \mathcal{T} is the fundamental tangle of $K_{a,b}$, then $K_{a,b}$ contains a \mathcal{T} -aligned $n \times n$ -grid-minor.*

Proof. Contracting a matching of size n^2 in $K_{a,b}$ and deleting the vertices not incident with this matching yields a K_{n^2} -minor, which has an $n \times n$ -grid subgraph, G' . Let \mathcal{T}' be the fundamental tangle of G' .

Suppose (A, B) is a $(n-1)$ -separation in $K_{a,b}$ and $A \cap E(G') \in \mathcal{T}'$. Then, $B \cap E(G') \notin \mathcal{T}'$ by tangle property (T4), so $B \cap E(G')$ contains the edge set of some row of G' . Each row of G' contains n vertices, so the number of vertices in $V(G')$ incident with an edge in B is at least n , and hence the number of vertices in $V(K_{a,b})$ incident with an edge in B is also at least n . Therefore, because (A, B) is an $(n-1)$ -separation in $K_{a,b}$, there is some vertex, v such that B contains each edge incident with v in $K_{a,b}$. Thus, $B \notin \mathcal{T}$, so $A \in \mathcal{T}$, so G' is \mathcal{T} -aligned. \square

Lemma 4.10.3. *For $r, \ell, m, n \in \mathbb{N}$ with $n \geq 3$, $r \geq 1$, $m \geq n^2$ and $\ell \geq n^2$, if W is an (r, ℓ, m) wheel with model $\mathcal{W} = (\mathbb{T}, \mathbb{M}, Z, \mathfrak{e})$ and \mathcal{T} is the fundamental tangle of \mathcal{W} , then W contains a \mathcal{T} -aligned $n \times n$ -grid-minor.*

Proof. Let G' be obtained from W by contracting the edges of each rim-tree T_i and deleting the edges of each rim-matching M_i . Thus, G' is isomorphic to $K_{\ell, m}$.

To see that G' is \mathcal{T} -aligned, let \mathcal{T}' be the fundamental tangle of G' and suppose (A, B) is a $(\ell - 1)$ -separation in W with $A \cap E(G') \in \mathcal{T}'$. It cannot be the case that each vertex in Z is incident with both an edge in A and an edge in B , so there exists $z \in Z$ such that each edge incident with z is in B . Therefore, B contains an edge in each rim-tree of \mathcal{W} , so $B \notin \mathcal{T}$. Hence, $A \in \mathcal{T}$, so G' is \mathcal{T} -aligned and, by Lemma 4.10.2, G' contains a \mathcal{T} -aligned $n \times n$ -grid-minor. \square

Lemma 4.10.4. *For $r, \ell, m, n \in \mathbb{N}$ with $n \geq 3$, $m \geq n^2$ and $r \geq n^3$, if W is a homogeneous (r, ℓ, m) wheel with model $\mathcal{W} = (\mathbb{T}, \mathbb{M}, Z, \mathfrak{e})$ and \mathcal{T} is the fundamental tangle of \mathcal{W} , then W contains a \mathcal{T} -aligned $n \times n$ -grid-minor.*

Proof. Because W is a homogenous wheel, there is some r -vertex tree T such that each rim-tree T_i in \mathcal{W} is a copy of T . For $v \in V(T)$ and $i \in \mathbb{Z}_n$, let v^i denote the copy of v in T_i ; for $e \in E(T)$ and $i \in \mathbb{Z}_n$, let e^i denote the copy of e in T_i .

Claim 4.10.4.1. *Either T contains at least n^2 leaves or T contains a path of length at least $n - 1$.*

Proof of Claim. Suppose T has no path of length $n - 1$. Let v_0 be leaf in T . For each leaf v in T , let P_v be the path in T between v_0 and v . By assumption, $|V(P_v)| \leq n$. Each vertex in T is in $V(P_v)$ for some leaf v , so

$$\bigcup_{v \in L(T)} V(P_v) = V(T).$$

Hence,

$$n|L(T)| \geq \sum_{v \in L(T)} |V(P_v)| \geq |V(T)| = r \geq n^3,$$

so $|L(T)| \geq n^2$, as desired. □ (Claim)

Claim 4.10.4.2. *If T has at least n^2 leaves, then W has a \mathcal{T} -aligned $n \times n$ -grid-minor.*

Proof of Claim. Let v_1, \dots, v_{n^2} be leaves of T , and let e_1, \dots, e_{n^2} be the edges incident with each of these leaves, respectively. For $i \in \{1, \dots, m - 1\}$ and $j \in \{1, \dots, n^2\}$, let $e'_{i,j}$ be the edge in M_i incident with v_j^i and v_j^{i+1} ; such an edge exists because W is a homogeneous wheel. Let G' be obtained from W by contracting, for each $i \in \{1, \dots, m - 1\}$,

$$(E(T_i) - \{e_1^i, \dots, e_{n^2}^i\}) \cup \{e'_{i,j} : j \in \{1, \dots, n^2\}\},$$

and deleting all remaining edges except

$$\{e_j^i : i \in \{1, \dots, m - 1\}, j \in \{1, \dots, n^2\}\}.$$

For $i \in \{1, \dots, m - 1\}$, let u_i be the vertex formed by the contraction of the edges $(E(T_i) - \{e_1^i, \dots, e_{n^2}^i\})$ (which is connected). For $j \in \{1, \dots, n^2\}$, let v'_j be the vertex formed by contracting $\{e'_{i,j} : i \in \{1, \dots, m - 1\}\}$. Note that, for $i \in \{1, \dots, m - 1\}$ and $j \in \{1, \dots, n^2\}$, e_j^i is incident with u_i and v'_j in G' . Hence, G' is isomorphic to $K_{n^2, m-1}$. Let \mathcal{T}' be the fundamental tangle of G' .

To see that G' is \mathcal{T} -aligned, let (A, B) be an $(n^2 - 1)$ -separation in W with $A \cap E(G') \in \mathcal{T}'$. It cannot be the case that each vertex in $\{v'_j : j \in \{1, \dots, n^2\}\}$ is incident with both

an edge in A and an edge in B , so there exists $j \in \{1, \dots, n^2\}$ such that each edge incident with v'_j is in B . Hence, for $i \in \{1, \dots, m-1\}$, $e_j^i \in B \cap E(T_i^+)$. By Lemma 4.3.2 and the fact that $m > n^2$, $B \notin \mathcal{T}$. Hence, $A \in \mathcal{T}$, so G' is \mathcal{T} -aligned and, by Lemma 4.10.2, G' contains a \mathcal{T} -aligned $n \times n$ -grid-minor. \square (Claim)

Claim 4.10.4.3. *If T has a path of length at least $n-1$, then W has an $n \times n$ -grid-minor.*

Proof of Claim. Let v_1, \dots, v_n be vertices of a path in T of length $n-1$. Because W is a homogeneous wheel, for $i \in \{1, \dots, m-1\}$, $j \in \{1, \dots, n\}$, the rim-matching M_i contains an edge $e'_{i,j}$ incident with vertices v_j^i and v_j^{i+1} . Therefore, the subgraph of W induced by the vertices $\{v_j^i : i, j \in \{1, \dots, n\}\}$ is an $n \times n$ -grid, G' . Let \mathcal{T}' be the fundamental tangle of G' . Label G' such that, for each $\{i \in \{1, \dots, n\}\}$, row i of the grid G' contains vertex set

$$\{v_j^i : j \in \{1, \dots, n\}\},$$

which is contained in rim-tree T_i .

To see that G' is \mathcal{T} -aligned, let (A, B) be an $(n-1)$ -separation in W with $A \cap E(G') \in \mathcal{T}'$. By the definition of the fundamental tangle of a grid, for each $i \in \{1, \dots, n\}$, row i of G' contains some edge in B ; therefore, the augmented tree T_i^+ contains an edge in B . Hence, because (A, B) is an $(n-1)$ -separation, there is some $i \in \{1, \dots, n\}$ such that each edge in the augmented tree T_i^+ is in B . Thus, $B \notin \mathcal{T}$, so $A \in \mathcal{T}$ and G' is \mathcal{T} -aligned. \square (Claim)

In both cases, W has a \mathcal{T} -aligned $n \times n$ -grid-minor. \square

Theorem 4.4.3. (Grid Theorem). *For each $n \in \mathbb{N}$, there exists some $N \in \mathbb{N}$ such that if G contains a tangle \mathcal{T} of order at least N , then G contains a \mathcal{T} -aligned $n \times n$ -grid-minor.*

Proof. Let $\theta = n^2 + n^3 - 1$, $m = n^2$, $\delta = f_{4.4.2}(\theta, m)$ and $N = 2^\delta \theta$. Suppose G is a graph with a tangle of order N . By Lemma 4.10.1, G contains a θ -tangle \mathcal{T} with covering-number at least δ . By Theorem 4.4.2, G contains a \mathcal{T} -aligned minor G' isomorphic to either $K_{\theta, m}$ -minor or a homogeneous (r, ℓ, m) -wheel with model \mathcal{W} , for some $r, \ell \in \mathbb{N}$ with $r \geq 1$ and $2r + \ell = \theta$.

Let \mathcal{T}' be either the fundamental tangle of G' , if G' is isomorphic to $K_{\theta, m}$, or the fundamental tangle of the model \mathcal{W} , if G' is isomorphic to an (r, ℓ, m) -wheel. Note that it suffices to show that G' contains a \mathcal{T}' -aligned $n \times n$ -grid-minor. If G' is isomorphic to $K_{\theta, m}$, then G' contains a \mathcal{T}' -aligned $n \times n$ -grid-minor by Lemma 4.10.2. If G' is isomorphic to an (r, ℓ, m) -wheel and $\ell \geq n^2$, then G' contains a \mathcal{T}' -aligned $n \times n$ -grid-minor by Lemma 4.10.3. Otherwise, G' is isomorphic to an (r, ℓ, m) -wheel and $r \geq n^3$, so G' contains a \mathcal{T}' -aligned $n \times n$ -grid-minor by Lemma 4.10.4. \square

Chapter 5

Conclusions and Further Work

This research was motivated by a desire to describe the structure of graphs with no K_6 -minor. This goal is far from complete, but we have found some promising tools for such a structure theorem.

Jørgensen’s Conjecture [14] (Conjecture 1.3.6), stating that 6-connected graphs with no K_6 -minor are apex, and Kawarabayashi, Norin, Thomas and Wollan’s theorem that Jørgensen’s Conjecture holds for sufficiently large 6-connected graphs, suggest that it should be possible to decompose graphs with no K_6 -minor along 5-separations into pieces that either have bounded size, or are “almost planar”, in the sense that deleting a small number of vertices yields a planar graph.

Guided by this intuition, we have shown in Chapters 2 and 3, that *any* graph has a tree-decomposition along 5-separations such that each node in the tree-decomposition either has small degree or displays a tangle of order 6 with large covering number. The piece of a graph corresponding to a small degree node in a width-5 tree-decomposition, can only have a small number of vertices. On the other hand, for nodes displaying 6-tangles with high covering-number, the tangle gives some local structure, but it is rather abstract, and far from constructive.

To better understand these 6-tangles with high covering-number, we described the highly structured families of minors that such tangles force in a graph. These minors are the (r, ℓ, n) -wheels described in Chapter 4, as well as $K_{6,n}$. With a little bit of Ramsey Theory, we showed that any 6-tangle with sufficiently high covering-number forces one of the minors depicted in Figure 5.1.

To push this work closer to a structure theorem for graphs with no K_6 -minor, it seems that we must understand graphs which contain one of the graphs depicted in Figure 5.1




Figure 5.1: The unavoidable-minors for graphs with 6-tangles of high covering-number.

as a minor, but do not contain K_6 as a minor. We suspect that such graphs are “almost planar” in some sense; perhaps, as we ventured to guess in Conjecture 1.4.8 these graphs are 4-apex.

Four of the graphs in Figure 5.1, (a), (b), (e) and (i) contain K_6 -minors, as shown in Figure 5.2; the remaining graphs are each planar or 1-apex.

Norin and Thomas [20] have described the minimal non-planar extensions of planar graphs, and developed some tools for understanding the minimal non- k -apex extensions of k -apex graphs. Unfortunately, these results do not immediately give an adequate description of the non-4-apex extensions of the six remaining graphs in Figure 5.1. However, the techniques developed by Norin and Thomas seem promising, and, with further development, might be used to complete this structure theorem.

It is also worth pointing out that the techniques developed in this thesis are quite general, with potential applications in proving structure theorems for graphs that do not contain certain other small graphs as minors. For example, one might be able to describe the structure of graphs with no Petersen Graph minor, or the structure of graphs with no minor in the Petersen Family, including K_6 , the Petersen Graph, and the five other graphs. Graphs with no minor in the Peterson Family are exactly the graphs that can be embedded in \mathbb{R}^3 in such a way that no two cycles in the graph are *linked*, in the sense of knot theory, so a structure theorem for graphs with no minor in the Petersen Family would be a structure theorem for linklessly embeddable graphs.

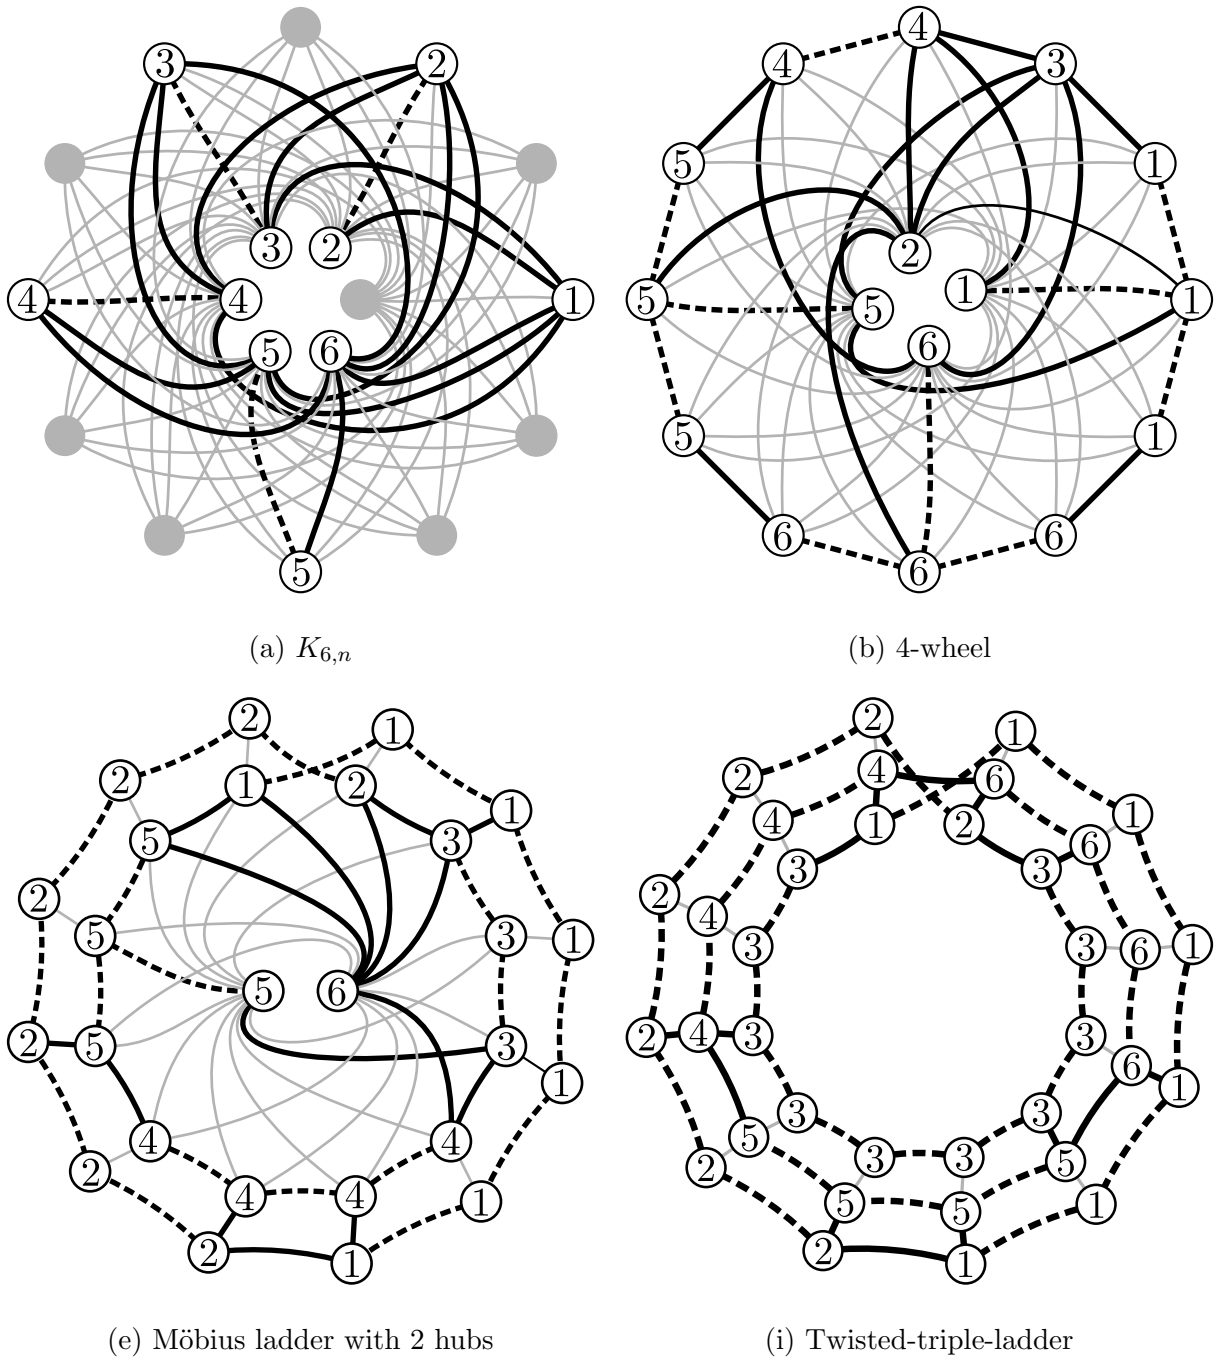


Figure 5.2: K_6 -minors in four of the graphs in Figure 5.1. The K_6 -minor is obtained by contracting the dashed edges and deleting the light grey edges.

Bibliography

- [1] Umberto Bertele and Francesco Brioschi. Nonserial dynamic programming. *Mathematics in Science and Engineering*, 91, 1972.
- [2] John H Conway and C. McA. Gordon. Knots and links in spatial graphs. *Journal of Graph Theory*, 7(4):445–453, 1983.
- [3] Reinhard Diestel. Graph theory (Graduate texts in mathematics). 2005.
- [4] Reinhard Diestel, Tommy R. Jensen, Konstantin Yu. Gorbunov, and Carsten Thomassen. Highly Connected Sets and the Excluded Grid Theorem. *Journal of Combinatorial Theory, Series B*, 75(1):61–73, January 1999.
- [5] Guoli Ding. A characterization of graphs with no octahedron minor. *Journal of Graph Theory*, 74(2):143–162, 2013.
- [6] Guoli Ding and Cheng Liu. A Chain Theorem for 3^+ -Connected Graphs. *SIAM Journal on Discrete Mathematics*, 26(1):102–113, 2012.
- [7] Guoli Ding and Cheng Liu. Excluding a small minor. *Discrete Applied Mathematics*, 161(3):355–368, 2013.
- [8] Richard J Duffin. Topology of series-parallel networks. *Journal of Mathematical Analysis and Applications*, 10(2):303–318, 1965.
- [9] Satoru Fujishige. Canonical decompositions of symmetric submodular systems. *Discrete Applied Mathematics*, 5(2):175–190, 1983.
- [10] Jim Geelen, Bert Gerards, Neil Robertson, and Geoff Whittle. Obstructions to branch-decomposition of matroids. *Journal of Combinatorial Theory. Series B*, 96:560–570, 2006.

- [11] Jim Geelen, Bert Gerards, and Geoff Whittle. Tangles, tree-decompositions and grids in matroids. *Journal of Combinatorial Theory, Series B*, 99(4):657–667, July 2009.
- [12] Jim Geelen and Kasper Kabell. The erdős–pósa property for matroid circuits. *Journal of Combinatorial Theory, Series B*, 99(2):407–419, 2009.
- [13] Rudolf Halin. S-functions for graphs. *Journal of Geometry*, 8(1-2):171–186, 1976.
- [14] Leif K Jørgensen. Contractions to K_8 . *Journal of Graph Theory*, 18(5):431–448, 1994.
- [15] Ken-ichi Kawarabayashi and Bojan Mohar. Some recent progress and applications in graph minor theory. *Graphs and Combinatorics*, 23(1):1–46, 2007.
- [16] Ken-ichi Kawarabayashi, Raiji Mukae, and Atsuhiko Nakamoto. K_6 -Minors in Triangulations on the Klein Bottle. *SIAM Journal on Discrete Mathematics*, 23(1):96–108, 2008.
- [17] Ken-ichi Kawarabayashi, Sergey Norin, Robin Thomas, and Paul Wollan. K_6 minors in large 6-connected graphs. *arXiv preprint arXiv:1203.2192*, 2012.
- [18] A K Kelmans. 3-connected graphs without essential 3-cuts and triangles. *Doklady Akademii Nauk Sssr*, 288(3):531–535, 1986.
- [19] John Maharry. A characterization of graphs with no cube minor. *Journal of Combinatorial Theory, Series B*, 80(2):179–201, 2000.
- [20] Sergey Norin and Robin Thomas. Non-planar extensions of subdivisions of planar graphs. *arXiv preprint arXiv:1402.1999*, 2014.
- [21] Bogdan Oporowski, James Oxley, and Robin Thomas. Typical subgraphs of 3- and 4-connected graphs. *Journal of Combinatorial Theory, Series B*, 57(2):239–257, 1993.
- [22] Neil Robertson. Excluding V_8 as a minor. *preprint*.
- [23] Neil Robertson and Paul Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41:92–114, 1986.
- [24] Neil Robertson and Paul Seymour. Graph minors. IV. Tree-width and well-quasi-ordering. *Journal of Combinatorial Theory, Series B*, 48:227–254, 1990.
- [25] Neil Robertson and Paul Seymour. Graph minors. X. Obstructions to tree-decomposition. *Journal of Combinatorial Theory, Series B*, 52(2):153–190, July 1991.

- [26] Neil Robertson and Paul Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory. Series B*, 89:43–76, 2003.
- [27] Neil Robertson, Paul Seymour, and Robin Thomas. Sachs Linkless Embedding Conjecture. *Journal of Combinatorial Theory, Series B*, 64(2):185–227, July 1995.
- [28] Neil Robertson, Paul D. Seymour, and Robin Thomas. Quickly excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 62:323–348, 1994.
- [29] Horst Sachs. On spatial representations of finite graphs. *Finite and infinite sets*, 1:649–662, 1984.
- [30] Robin Thomas. Recent excluded minor theorems for graphs. *London Mathematical Society Lecture Note Series*, pages 201–222, 1999.
- [31] Klaus Wagner. Über eine Eigenschaft der ebenen Komplexe. *Mathematische Annalen*, 114(1):570–590, 1937.
- [32] Peter Whalen. *Pfaffian Orientations, Flat Embeddings, and Steinberg’s Conjecture*. PhD thesis, Georgia Institute of Technology, 2014.

Notation

(X, Y) -path a path between vertex sets X and Y , page 97

$\kappa_\lambda(A, B)$ connectivity between A and B , page 39

θ -separation a separation of order at most θ , page 7

θ -tangle tangle of order θ , page 26

θ -tangle tangle of order θ , page 41

θ -tree-width the minim node-width of the θ -tree-decomposition of a graph, page 23

H -minor-free a graph with no minor isomorphic to H , page 2

$I(T)$ the set of node-edge incidences in a tree T , page 58

k -sum a clique-sum of order k , page 2

$L(T)$ the set of leaves of the tree T , page 45

T_i^+ augmented tree, page 88

λ a connectivity function, page 38

λ_G the connectivity function derived from a graph, page 39

Index

- H -minor, 2
- ∞ -robust, 65
- θ -branch-degree, 23
- θ -complete, 73
- θ -connected set, 30, 82
- θ -elementary, 60
- θ -tree-decomposition, 23
- θ -tree-width, 23
- k -apex, 5
- k -sum, 2
- (BDPTT), 72
- (CMC), 73
- (DC1), 70
- (DC2), 70
- (FTT), 60
- (INIT), 62
- (M1), 49
- (M2), 49
- (M3), 49
- (MC1), 73
- (MC2), 73
- (MC3), 73
- (MDC), 71
- (N1), 97
- (N2), 97
- (N3), 97
- (N4), 98
- (PTT), 65
- (T1), 26, 41
- (T2), 26, 41
- (T3), 26, 41
- (T4'), 41
- (T4), 27, 41
- (T5), 27, 41
- (TB1), 69
- (TB2), 69
- (TIGHT), 73
- (TT1), 60, 64
- (TT2), 60, 65
- 0-sum, 2
- 1-sum, 2
- 2-wheel, 80
- aligned
 - necklace, 98
- aligned $K_{\theta,n}$ -minor, 94
- aligned wheel-minor, 94
- apex, 9
 - k -apex, 5
- attachment vertex, 102
- augmented tree, 88
- bead, 97
 - weak, 98
- bounded-degree, 72
- bounded-degree tree of θ -tangles, 62
- branch-degree
 - θ -branch-degree, 23, 41
- branch-width, 22
- bridge, 102
 - attachment vertex, 102

- chain, 114
- clique-sum, 2
- complete, 73
- connected
 - θ -connected, 7
 - θ -connected set, 30, 52
 - internally 4-connected, 8
- connectivity, 39
- connectivity function, 38
- connectivity system, 38
 - graphic, 39
- covered
 - m -covered, 103
- covering-number, 28, 43
- cross, 18
- cyclic-ladder, 80
- cylindrical grid, 43
- displays, 24, 25, 29
 - tangle in connectivity system, 58
- distinguishing covering, 70
 - minimum, 71
- distinguishing separation, 59
 - $\hat{\lambda}$ -minimum, 69
 - minimum-order, 59
- dominates, 39
- doublecross, 12
- edge-bag, 16
- edge-contraction, 1
- edge-width, 22, 41
- elementary
 - θ -elementary, 60
- embed, 5
- exact structure theorem, 3
- extended connectivity system, 43
- full tree of θ -tangles, 60
- fundamental tangle
 - of $K_{\theta,n}$, 91
 - of a grid, 96
 - of a wheel model, 88
- generalized wheel, 83
- graphic connectivity system, 39
- grid, 30, 82
- Grid Theorem, 31
- ground set, 38, 49
- hamburger graph, 13
- homogeneous wheel, 84
- hose graph, 12
- hub, 86
 - of a family of subgraphs, 102
 - of a necklace, 98
- hub (of (r, ℓ, n) -wheel), 83
- hyperplane, 54
- incidence, 40
- induced, 42
- initial, 62
- internally 4-connected, 8
- jump, 111
- jump-free, 111
- Jørgensen's Conjecture, 9
- leaf, 40
- leaf-incidence, 40
- length (of (r, ℓ, n) -wheel), 84
- linked necklace, 99
- linkless embedding, 12
- matroid, 49
- maximal tangle, 60
- minimum-order distinguishing separation, 59
- minor, 1
 - H -minor-free, 2
- model of a wheel, 88

- monotone covering, 73
- monotonic, 49
- Möbius-ladder, 80
- necklace, 97
- node-bag, 16
- node-width, 22
- order
 - of a connectivity system separation, 38
 - of a separation, 7
- part, 18
- partial tree of θ -tangles, 64
 - bounded-degree, 72
- partial tree-decomposition, 45, 63
 - tight, 73
- path
 - (X, Y) -path, 97
 - between X and Y , 97
- planar-plus-triangle, 11
- present, 18
- presented
 - separation, 41, 45
 - set, 41, 45
- rank function, 49
- restriction, 54
- rim-matching, 83
- rim-tree, 83
- robust, 65
 - ∞ -robust, 65
- row, 30, 42, 96
- separation
 - θ -separation, 7
 - in a connectivity system, 38
 - in a graph, 7
- separation presented, 41, 45
- set presented, 41, 45
- small, 27
- spoke, 83
- submodular, 38
- supported necklace, 99
- symmetric, 38
- tangle, 26, 41
 - truncation, 136
- tangle matroid, 82
- tangle-matroid, 50
- tiebreaker, 69
- tight, 73
- transversal, 93
- tree of θ -tangles, 60
 - bounded-degree, 62
 - full, 60
 - partial, 64
- tree-decomposition, 16
 - θ -tree-decomposition, 23
 - connectivity system, 40
 - partial, 63
- tree-width, 22
 - θ -tree-width, 23
- truncation of a tangle, 136
- twisted-triple-ladder, 88
- uniform matroid, 54
- Wagner's Theorem, 3
- weak bead, 98
- wheel, 80
 - (r, ℓ, n) -wheel, 32, 83
 - homogeneous, 84
- width, 45, 73