All graphs have tree-decompositions displaying their topological ends

Carmesin, Johannes

DOI:
10.1007/s00493-018-3572-0

All graphs have tree-decompositions displaying their topological ends

Carmesin, Johannes

All graphs have tree-decompositions displaying their topological ends

Johannes Carmesin

University of Cambridge

Abstract

We show that every connected graph has a spanning tree that displays all its topological ends. This proves a 1964 conjecture of Halin in corrected form, and settles a problem of Diestel from 1992.

1 Introduction

In 1931, Freudenthal introduced a notion of *ends* for second countable Hausdorff spaces [20], and in particular for locally finite graphs [21]. Independently, in 1964, Halin [23] introduced a notion of *ends* for graphs, taking his cue directly from Carathéodory's *Primenden* of simply connected regions of the complex plane [4]. For locally finite graphs these two notions of ends agree.

For graphs that are not locally finite, Freudenthal's topological definition still makes sense, and gave rise to the notion of *topological ends* of arbitrary graphs [17]. In general, this no longer agrees with Halin's notion of ends, although it does for trees.

Halin [23] conjectured that the end structure of every connected graph can be displayed by the ends of a suitable spanning tree of that graph. He proved this for countable graphs. Halin's conjecture was finally disproved in the 1990s by Seymour and Thomas [27], and independently by Thomassen [30].

In this paper we shall prove Halin's conjecture in amended form, based on the topological notion of ends rather than Halin's own graph-theoretical notion. We shall obtain it as a corollary of the following theorem, which proves a conjecture of Diestel [13] of 1992 (again, in amended form):

Theorem 1. *Every graph has a tree-decomposition $(T, \mathcal{V})$ of finite adhesion such that the ends of T define precisely the topological ends of G.*
See Section 2 for definitions.

The tree-decompositions constructed for the proof of Theorem 1 have several further applications. In [6] we use them to answer the question to what extent the ends of a graph - now in Halin's sense - have a tree-like structure at all. In [8], we apply Theorem 1 to show that the topological cycles of any graph together with its topological ends induce a matroid. We remark that although the existence of a tree-decomposition as in Theorem 1 for an arbitrarily subset of the vertex-ends in place of the topological ends implies the existence of a suitable spanning tree in Halin's sense for that subset by Remark 6.14, the converse is not true, see Example 3.1.

This paper is organised as follows. In Section 2 we explain the problems of Diestel and Halin in detail, after having given some basic definitions. In Section 3 we continue with examples related to these problems. Section 4 only contains material that is relevant for Section 5 in which we prove that every graph has a nested set of separations distinguishing the vertex-ends efficiently. In Section 6, we use this theorem to prove Theorem 1. Then we deduce Halin's amended conjecture. Finally, Section 7 contains concluding remarks.

2 Definitions

Throughout, notation and terminology for graphs are that of [14]. And G always denotes a graph.

A *vertex-end* in a graph G is an equivalence class of rays (one-way infinite paths), where two rays are equivalent if they cannot be separated in G by removing finitely many vertices. Put another way, this equivalence relation is the transitive closure of the relation relating two rays if they intersect infinitely often.

Example 2.1. The vertex-ends of rooted trees are (in bijection with) the rays starting at the root; of course vertex-ends do not depend on the choice of a root.

Let X be a locally connected Hausdorff space. Given a subset $Y \subseteq X$, we write $\overline{Y}$ for the closure of Y, and $F(Y) := \overline{Y} \cap \overline{X \setminus Y}$ for its frontier. In order to define the topological ends of X, we consider infinite sequences $U_1 \supseteq U_2 \supseteq ...$ of non-empty connected open subsets of X such that each $F(U_i)$ is compact and $\bigcap_{i\geq 1} \overline{U}_i = \emptyset$. We say that two such sequences $U_1 \supseteq U_2 \supseteq ...$ and $U_1' \supseteq U_2' \supseteq ...$ are *equivalent* if for every i there is some j with $U_i \supseteq U_j'$. This relation is transitive and symmetric [20, Satz 2]. The equivalence classes of those sequences are the *topological ends* of X [17, 20, 26].

For the simplicial complex of a graph G, Diestel and Kühn described the topological ends combinatorically: a vertex *dominates* a vertex-end ω if for some (equivalently: every) ray R belonging to ω there is an infinite fan of v-R-paths that are vertex-disjoint except at v. In [17], they proved that the topological ends are given by the undominated vertex-ends. Hence in this paper, we take this as our definition of *topological end of* G.

Example 2.2. For locally finite graphs the notions of vertex-ends and topological ends agree.

Example 2.3. For trees the notions of vertex-ends and topological ends agree. Hence we just call the vertex-ends of trees *ends*.

For us, a *separation* is an (ordered) pair (A, B) of vertex sets A and B such that no edge has an endvertex in $A \setminus B$ and the other endvertex in $B \setminus A$. The set $A \cap B$ is called the *separator* of (A, B). The size of the separator is the *order* of (A, B). The sets A and B are called the *sides* of the separation. The *reverse* of the separation (A, B) is the separation (B, A).

Given two separations (A, B) and (C, D), we write $(A, B) \leq (C, D)$ if $A \subseteq C$ and $D \subseteq B$. These separations are *nested* if $(A, B) \leq (C, D)$ or one of the other three possibilities obtained by replacing (A, B) or (C, D) by their reverse. Formally, (A, B) and (C, D) are nested if $(A, B) \leq (C, D)$, $(B, A) \leq (C, D)$, $(A, B) \leq (D, C)$ or $(B, A) \leq (D, C)$.

Remark 2.4. Most separations of interest are 'proper', see below. By Observation 2.5, proper separations (A, B) and (C, D) satisfy $(A, B) \leq (C, D)$ already if $A \subseteq C$. In this sense our definition of nestedness corresponds to the notion of nestedness for sets.

A separation (A, B) is *proper* if every vertex in the separator $A \cap B$ has a neighbour in $A \setminus B$ and $B \setminus A$.

Observation 2.5. *For proper separations* (A, B) *and* (C, D) *the following are equivalent.*

1. $(A, B) \leq (C, D)$*;*

2. $A \subseteq C$*;*

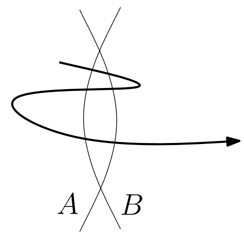

Figure 1: Every ray traverses the finite separator $A \cap B$ finitely often and then is eventually included in one of the sides A or B.

3. $A \setminus B \subseteq C \setminus D$.

Proof. As (A, B) is proper, the side A is determined by the set $A \setminus B$; indeed, it is $A \setminus B$ together with its neighbourhood. Conversely, also the side A determines the set $A \setminus B$: this set consists of those vertices of A that have all their neighbours in A. So (2) and (3) are equivalent.

Clearly (1) implies (2). Now conversely assume that $A \subseteq C$. By the above it suffices to show that $D \setminus C$ is included in $B \setminus A$. In other words: $G \setminus C$ is included in $G \setminus A$. This follows from $A \subseteq C$. □

A vertex-end ω *lives* in a side B of a separation (A, B) of finite order if the side B includes a ray belonging to ω. In this case B includes a subray of every ray belonging to ω, see Figure 1. A separation (A, B) of finite order *distinguishes* two vertex-ends ω and μ if one of them lives in the side A and the other lives in the side B. It distinguishes them *efficiently* if (A, B) has minimal order amongst all separations distinguishing ω and μ.

A *tree-decomposition* of a graph G consists of a tree T together with a family of subgraphs[1] $(P_t | t \in V(T))$ of G such that every vertex and edge of G is in at least one of these subgraphs, and such that if v is a vertex of both P_t and P_w, then it is a vertex of each P_u, where u lies on the v-w-path in T. We call the subgraphs P_t the *parts* of the tree-decomposition. The *adhesion* of a tree-decomposition is finite if adjacent parts intersect only finitely. Given an edge tu of T, we denote by T_t the subtree of $T - tu$ that contains t. Given a directed edge tu of T, the *separation corresponding to tu* is the separation (A_t, A_u), where A_i is the union of all parts P_x with $x \in T_i$ for $i = u, t$.

In [2, 25, 29], tree-decompositions of finite adhesion are used to study the

[1]We denote the vertex set of a graph G by $V(G)$.

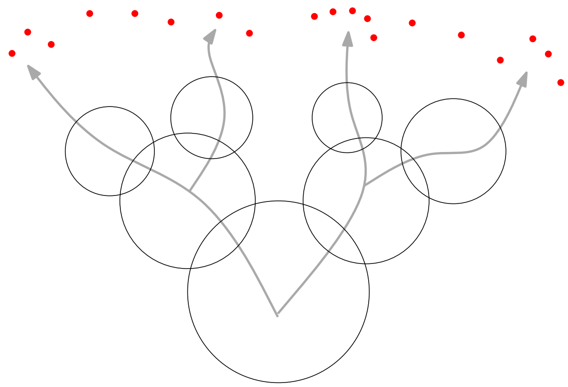

Figure 2: The ends of the decomposition tree - this tree is indicated in grey - define precisely the vertex-ends of the graph indicated by dots. Other vertex-ends living in parts are not drawn.

structure of infinite graphs. In [13, Problem 4.3], Diestel wanted to know whether every graph G has a tree-decomposition $(T, P_t | t \in V(T))$ of finite adhesion that somehow encodes the structure of the graph with its ends.

Let us be more precise. Given a vertex-end ω, we take $O(\omega)$ to consist of those oriented edges tu of T such that ω lives in its corresponding separation. Note that $O(\omega)$ contains precisely one of the two directions tu and ut of each edge of the tree. Furthermore this orientation $O(\omega)$ of T points towards a node of T or to an end of T. We say that ω *lives* in the part for that node or that end, respectively.

A vertex-end ω is *thin* if every set of vertex-disjoint rays belonging to ω is finite; otherwise ω is *thick*. Diestel asked whether every graph has a tree-decomposition $(T, P_t | t \in V(T))$ of finite adhesion such that different thick vertex-ends live in different parts and such that the ends of T *define precisely* the thin vertex-ends; here the ends of T *define precisely* a set Ψ of vertex-ends of G if in every end of T there lives a unique vertex-end and it is in Ψ and conversely every vertex-end in Ψ lives in some end of T, see Figure 2.

Unfortunately, that is not true; in Example 3.1, we construct a graph such that each of its tree-decompositions of finite adhesion has a part in which two (thick) vertex-ends live. In Example 3.3 we refine that construction by constructing a graph such that there live two thin vertex-ends in some part of every such a tree-decomposition.

Hence there remains the open question whether there is a natural subclass of the vertex-ends (similar to the class of thin vertex-ends) such that every graph has a tree-decomposition of finite adhesion such that the ends of its decomposition tree define precisely the vertex-ends in that subclass. Another question that arises in this context is: what is the largest possible natural class of vertex-ends such that every graph has a tree-decomposition distinguishing the vertex-ends in that class? Theorem 1 above answers the first question affirmatively. In Section 7, we show how Theorem 1 can be used to obtain a satisfying answer to the second question.

It is impossible to construct a tree-decomposition as in Theorem 1 with the additional property that for any two topological ends ω and μ, there is a separation corresponding to an edge of the tree that separates ω and μ efficiently, see Example 3.7.

A recent development in the theory of infinite graphs seeks to extend theorems about finite graphs and their cycles to infinite graphs and the topological circles formed with their ends, see for example [1, 3, 18, 19, 22, 28], and [12] for a survey. We expect that Theorem 1 has further applications in this direction aside from the one mentioned in the introduction.

A rooted spanning tree T of a graph G is *end-faithful* for a set Ψ of vertex-ends if each vertex-end $\omega \in \Psi$ is uniquely represented by T in the sense that T contains a unique ray belonging to ω and starting at the root. For example, every normal spanning tree is end-faithful for all vertex-ends. Halin conjectured that every connected graph has an end-faithful tree for all vertex-ends. At the end of Section 6, we show that Theorem 1 implies the following nontrivial weakening of this disproved conjecture:

Corollary 2.6. *Every connected graph has an end-faithful spanning tree for the topological ends.*

One might ask whether it is possible to construct an end-faithful spanning tree for the topological ends with the additional property that it does not include any ray to any other vertex-end. However, this is not possible in general. Indeed, Seymour and Thomas constructed a graph G with no topological end that does not have a rayless spanning tree [27].

3 Example section

Example 3.1. In this example we give two constructions of graphs that have no tree-decompositions of finite adhesion that distinguish all vertex-ends. These constructions motivate the construction of Example 3.3, where

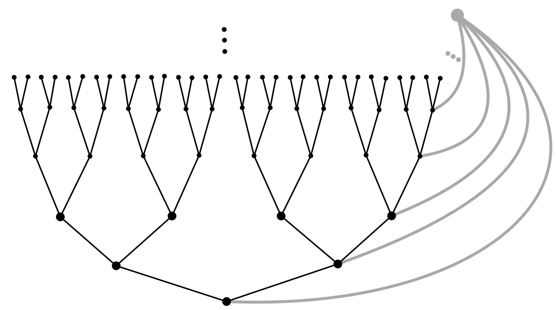

Figure 3: The binary tree is indicated in black. In grey we indicated the addition of a top along the right most ray.

we can construct such a graph not only for the class of vertex-ends but for the finer class of thin vertex-ends.

The simplest example of a graph with no such tree-decomposition for the vertex-ends known to the author is the (infinite) binary tree with tops, see Figure 3; this graph is obtained from the binary tree T_2 by adding one new vertex for every ray starting at the root. This new vertex is adjacent to all vertices on that ray. We call these new vertices the *tops*. We omit the proof that this graph has no tree-decompositions of finite adhesion that distinguishes all vertex-ends[2].

A slightly more complicated example is obtained from the regular tree T_ω with countably infinite degree by adding fat tops; here adding *fat tops* means that at each ray of T_ω starting at the root, we attach uncountably many, say $\aleph_1$, tops (that is new vertices adjacent to all vertices on the ray).

We sketch the proof that T_ω with fat tops has no tree-decompositions of finite adhesion that distinguishes all vertex-ends. First one checks that the vertex-ends of T_ω with fat tops are the ends of T_ω (this proof is similar to Lemma 3.4 below). The vertex-ends of T_ω with fat tops, however, are *fat*, that is, they are dominated by uncountably many vertices. The key observation is the following.

Lemma 3.2. *Let H be any graph with a tree-decomposition $(T, \mathcal{V})$ of finite adhesion. Then no fat vertex-end of H lives in an end of T.*

Proof. Vertex-ends living in ends μ of T can only be dominated by those vertices that eventually are in the separators corresponding to the edges on some ray in μ. Since the tree-decomposition has finite adhesion, there can

[2]A proof can be found in an earlier version of this paper [5]

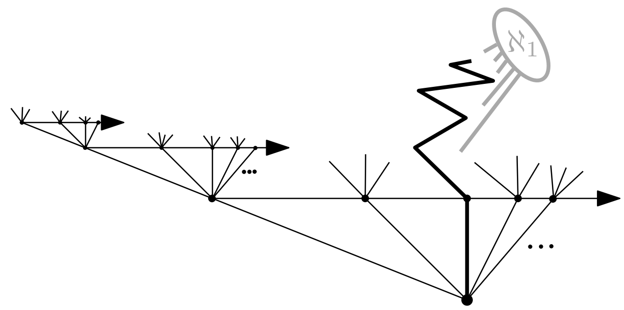

Figure 4: The graph G' is indicated in black. We indicated in grey the addition of fat tops at the highlighted ray. We obtain the graph G from the graph G' by adding these fat tops at all rays starting at the root.

only be countably many such vertices. So vertex-ends living in ends of the decomposition tree cannot be fat. □

In the final step one assumes that some tree-decomposition of finite adhesion distinguishes all vertex-ends. Since the graph T_ω is countable, it can only have countably many separators. A finite separator of T_ω with fat tops separates the same vertex-ends as their restriction to T_ω does. This essentially means[3] that the decomposition tree has only countably many edges. So it can only have countably many nodes. Since there are uncountably many vertex-ends, two of them have to live in the same part as they cannot live in an end of the decomposition tree by Lemma 3.2.

We remark that this proof also works for any graph obtained from T_ω by attaching some $\aleph_1$ fat tops at T_ω. So there is a counterexample against the statement that every graph has a tree-decomposition of finite adhesion distinguishing its vertex-ends of cardinality $\aleph_1$ – which is independent of the Continuum Hypothesis.

Example 3.3. In this example we construct a graph G such that each of its tree-decomposition of finite adhesion cannot distinguish all thin vertex-ends.

We start the construction with the regular tree T_ω of countably infinite degree. For each vertex of T_ω, we add a ray through its neighbours in the next level. Call the resulting graph G', see Figure 4. The vertex-ends of G'

[3]By contracting edges of the decomposition tree if necessary, we may assume that any two separations corresponding to edges of the decomposition tree distinguish different sets of vertex-ends. So no two such separations can have the same restriction to T_ω. Hence we may assume that the decomposition tree has only have countably many edges.

are those of T_ω together with one vertex-end for every newly added ray.

We obtain G from G' by adding for every ray of T_ω starting at the root a clique of uncountable cardinality $\aleph_1$ that is complete to that ray.

Lemma 3.4. *The vertex-ends of G' are (in bijection with) the vertex-ends of G.*

Proof. Every ray of G is equivalent to a ray of G'. Conversely any two vertex-ends of G' can be separated by a path of T_ω starting at the root. This path still separates rays belonging to these vertex-ends of G' in G. Hence G and G' have the same vertex-ends.

□

The thin vertex-ends of G are those vertex-ends of G' coming from newly added rays; indeed, if we remove the finite path of T_ω below such a newly added ray, all vertices on that ray become cut-vertices. All other vertex-ends are each dominated by uncountably many vertices, that is, they are fat.

We use the vertices of T_ω to refer to the thin vertex-ends. More precisely, we say that the vertex-end *sitting above* a vertex v is the one to which the ray in the upward neighbourhood of v belongs.

Suppose for a contradiction that the graph G has a tree-decomposition $(T, P_t | t \in V(T))$ of finite adhesion that distinguishes all its thin vertex-ends. First we show the following.

Lemma 3.5. *There is a ray R of T such that a fat vertex-end of G lives in the end to which R belongs.*

Proof. Our aim is to construct a sequence $(v_n | n \in \mathbb{N})$ of vertices that lie on a ray of the tree T_ω starting at the root together with a sequence $((A_n, B_n) | n \in \mathbb{N}^*)$ of separations corresponding to edges of the decomposition tree such that $(A_n, B_n) \leq (A_{n+1}, B_{n+1})$ and v_n is contained in $B_n \setminus A_n$.

We start the construction by picking an arbitrary separation (C, D) corresponding to an edge of the decomposition tree such that it distinguishes two thin vertex-ends. We pick for v_0 the root of the tree T_ω. By replacing the separation (C, D) by its reverse (D, C) if necessary, we may assume that the thin vertex-end sitting above v_0 lives in the side C. We let $(A_1, B_1) = (C, D)$. Let μ be a thin vertex-end living in B_1 and let u be the vertex of T_ω above which μ sits. The vertex u must be contained in the side B_1 and have all but finitely many of its upward-neighbours in the side B_1. Since the separator $A_1 \cap B_1$ is finite, the vertex u has an upward-neighbour v_1 in the rooted tree T_ω that is contained in $B_1 \setminus A_1$. We let P_1 be the unique path included in the tree T_ω from the vertex v_0 to the vertex v_1.

Now assume that we already constructed a path P_n and a separation (A_n, B_n) corresponding to an edge of the decomposition tree such that the last vertex v_n of P_n is contained in $B_n \setminus A_n$ and is a vertex of T_ω. Next we construct the path P_{n+1} and the separation (A_{n+1}, B_{n+1}). As the separator $A_n \cap B_n$ is finite, the vertex v_n has two upward-neighbours u and u' in the rooted tree T_ω contained in $B_n \setminus A_n$. By assumption there is a separation (C, D) corresponding to an edge of the decomposition tree such that the thin vertex-ends sitting above u and u' are distinguished by (C, D).

Sublemma 3.6. *The separation (A_n, B_n) is $\leq$ to the separation (C, D) or its reverse (D, C).*

Proof. This is a simple consequence of the fact that the separations (C, D) and (A_n, B_n) are nested as separations corresponding to edges of a decomposition tree of the same tree-decomposition.

The sides C and B_n both contain all but finitely many vertices of every ray belonging to the vertex-end sitting above the vertex u. Hence the intersection $C \cap B_n$ is infinite. Similarly, we conclude that the intersection $D \cap B_n$ is infinite. As the separator $C \cap D$ is finite, the side B_n cannot be included in one of the sides C or D. Hence as the separations (C, D) and (A_n, B_n) are nested, it must be that the separation (A_n, B_n) is $\leq$ to the separation (C, D) or its reverse (D, C). $\square$

By replacing the separation (C, D) by its reverse (D, C) if necessary we may assume by Sublemma 3.6 that $(A_n, B_n) \leq (C, D)$. We let $(A_{n+1}, B_{n+1}) = (C, D)$. Since the separator $A_{n+1} \cap B_{n+1}$ is finite and the thin vertex-end sitting above u' lives in B_{n+1}, the vertex u' has an upward-neighbour v_{n+1} in the rooted tree T_ω contained in $B_{n+1} \setminus A_{n+1}$. We obtain the path P_{n+1} from P_n by adding the unique path included in T_ω from the vertex v_n to the vertex v_{n+1}.

This completes the construction of the paths P_n and the separations (A_n, B_n). Hence by recursion, there is a sequence $(v_n | n \in \mathbb{N})$ of vertices that lie on a ray S of the tree T_ω starting at the root together with a sequence $((A_n, B_n) | n \in \mathbb{N}^*)$ of separations corresponding to edges of the decomposition tree such that $(A_n, B_n) \leq (A_{n+1}, B_{n+1})$ and v_n is contained in $B_n \setminus A_n$. The vertex-end μ to which the ray S belongs is an end of the tree T_ω; and thus is fat in the graph G. Since the ray S contains infinitely many vertices of all sides B_n, its vertex-end μ lives in all sides B_n. The edges corresponding to the separations (A_n, B_n) lie on a ray R of the decomposition tree; and the vertex-end μ lives in the end of R. This completes the proof. $\square$

Lemma 3.5 contradicts Lemma 3.2. This is the desired contradiction. Hence G has no tree-decomposition of finite adhesion that distinguishes all its thin vertex-ends.

Example 3.7. In this example, we construct a graph G such that for any of its tree-decompositions $(T, P_t | t \in V(T))$ there are two topological ends such that no separation corresponding to an edge of T distinguishes them efficiently[4].

We start the construction with the (cartesian) product[5] W of a ray with the path of five vertices, see Figure 5. By $P[n]$, we denote a graph that has

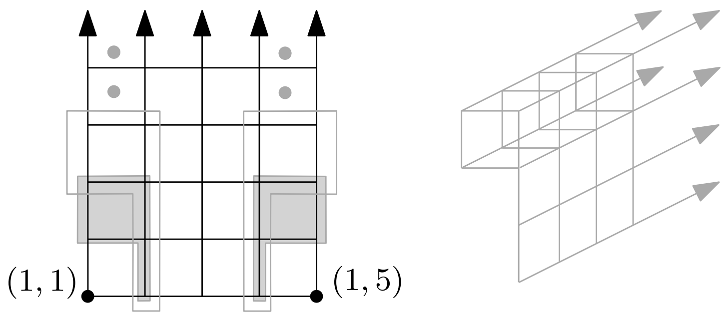

Figure 5: The construction of the graph G. It is obtained from the graph depicted on the left by attaching on each set of vertices surrounded by a 'P'-shaped box a graph like the one on the right of the appropriate size.

the shape of a 'P'. More precisely, it is obtained from a path of n vertices by adding an edge such that one endvertex of the edge is joined to the last vertex of the path and the other endvertex to the second but last. By H_n we denote the product of a ray with the graph $P[n]$. We obtain G from W by for each $n \geq 3$ attaching two copies of H_n as follows. We attach these new graphs H_n on copies of $P[n]$. The first copy is that containing the initial path of the ray of length n times the second vertex of the five-path together with the edge whose endvertices are the n-th and $(n-1)$-st vertex of the ray times the first vertex of the five-path. The second copy is that containing the initial path of the ray of length n times the forth vertex of the five-path together with the edge whose endvertices are the n-th and $(n-1)$-st

[4]Topological ends are examples of vertex-ends. In this sense the term 'distinguishes efficiently' is defined.

[5]Given two graphs G and H, by $G \times H$, we denote the graph with vertex set $V(G) \times V(H)$ where we join two vertices (g, h) and (g', h') by an edge if both $g = g'$ and $hh' \in E(G)$ or both $h = h'$ and $gg' \in E(G)$.

vertex of the ray times the fifth vertex of the five-path. In Figure 5 these attachment sets are surrounded by grey 'P'-shaped boxes. This completes the construction of G.

The vertex-ends of the attached graphs H_n are clearly topological. The graph G has the property that although we attach the graphs H_n at a copy of $P[n]$, the vertex-end of a new graph H_n can be separated from the vertex-end of the other copy of H_n by a separator properly contained in the attachment set $P[n]$; namely just those vertices in the attachment set that in W have a neighbourhood in the infinite component of W without the attachment set. The set of these vertices has the shape of an 'L' turned around and consists of $n+1$ vertices. We denote these separators by S_n^1 and S_n^2, depending on whether they are contained in the first or second attachment set $P[n]$, respectively.

It is straightforward to check that any separation separating the two vertex-ends of the two attached copies of H_n efficiently has the separating set S_n^1 or S_n^2.

Suppose for a contradiction that G has a tree-decomposition $(T, P_t | t \in V(T))$ that separates any two topological ends efficiently. Then infinitely many of its separations must have separating sets of the form S_n^1 or S_n^2. By symmetry we may assume that there are infinitely many of the form S_n^1.

By $(1,1)$ we denote the vertex of G that is the product of the first vertex of the five-path and first vertex of the ray, see Figure 5. Similarly, by $(1,5)$ we denote the vertex of G that is the product of the last vertex of five-path and first vertex of the ray. Let P_a be a part of the tree-decomposition that contains $(1,1)$ and similarly let P_b be a part of the tree-decomposition that contains $(1,5)$. The edges corresponding to the separations with separators of the form S_n^1 separate in T the vertex a from the vertex b; that is, they lie on the unique a-b-path. Since this path is finite, we derive the desired contradiction. Thus G has no tree-decomposition $(T, P_t | t \in V(T))$ such that for any two topological ends there is a separation corresponding to an edge of T distinguishes them efficiently.

We remark that all topological ends of G are thin vertex-ends and so this construction also shows that thin vertex-ends cannot always be distinguished efficiently.

4 Separations and tangles

In this section, we define tangles and related concepts and prove some intermediate lemmas that we will apply in Section 5.

4.1 Tangles

Tangles are a central concept in Graph Minor Theory that describe highly connected substructures of a graph such as complete subgraphs or grid minors. They do not explicitly describe these substructures. Instead, for every low order separation they point towards a side, where that substructure 'lives'. This side is called the *big* side and the other side of the separation is the *small* side. These assignments have to satisfy certain rules such as sides including big sides are big.

Formally, a *tangle of order* $k+1$ assigns to each separation of order[6] at most k a big side. The other side is called small. These assignments satisfy the following properties:

1. three small sides A_1, A_2, A_3 cannot cover all edges, in formulas: $G \neq G[A_1] \cup G[A_2] \cup G[A_3]$;
2. if X is a set of at most k vertices, there is a component C of $G - X$ such that $C \cup X$ is the big side of the separation $(C \cup X, G \setminus C)$.

From the first property it follows that if (A, B) is a separation of order at most k and $A \subseteq B$, then A is small and B is big in any tangle of order $k+1$. In particular, the empty set $\emptyset$ is the small side of $(\emptyset, G)$. Furthermore every separation of order at most k has precisely one big side in a tangle of order $k+1$ by the first property. And a side including a big side cannot be small. Thus if a side is a big side of some separation, it must be the big side of any separation it is a side of. Thus we shall say things like 'A is big' without specifying a separation (A, B) of which A is the big side.

Remark 4.1. In the standard definition of tangles for finite graphs (or more generally for locally finite graphs), the second property is omitted. The reason is that for finite graphs there is a simple well-known argument that it follows from the first. This argument relies on an induction on the number of components of $G - X$ and this implication no longer holds for quite simple infinite graphs like the infinite star (in fact without the second conditions non-principle ultra-filters on the leaves of the infinite star would give rise to a tangle of infinite order). It is not the scope of this paper to analyse such objects.[7] Hence we require this second condition.

We refer to this second condition as the *component property*.

[6] We follow the convention that we allow $k+1$ to be infinite. In that case we just replace 'of order at most k' by 'of finite order' in the above definition.

[7] Tangles without this second condition are studied in [16] by Diestel.

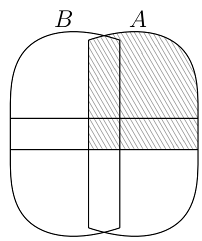

Figure 6: The corner diagram for the two separations (A, B) and (C, D). The separation (A, B) separates vertically, while (C, D) separates horizontally. The corner $A \cap C$ is shaded in grey. The middle region $A \cap B \cap C \cap D$ is called the *center*. The four other regions 'linking the corners' are called the *links*. Formally, they are $(A \cap B) \setminus C$, $(A \cap B) \setminus D$, $(C \cap D) \setminus A$ and $(C \cap D) \setminus B$.

In this paper we are mostly interested in the following examples of tangles.

Example 4.2. Each vertex-end ω induces a tangle; indeed, for a finite order separation (A, B) we define A to be big in this tangle if ω lives in A. It is straightforward to check that this defines a tangle of infinite order.

Given two separations (A, B) and (C, D), the separation $(A \cap C, B \cup D)$ is called the *corner separation* at the *corner* $A \cap C$, see Figure 6. In Figure 6 the separator of $(A \cap C, B \cup D)$ has the shape of an 'L'. Hence we denote this separator by $L(A, C)$; formally, $L(A, C)$ is the intersection of $A \cap C$ and $B \cup D$. The pair consisting of (A, B) and (C, D) has three more corner separations, corresponding to the corners of Figure 6. These are $(A \cap D, B \cup C)$, $(B \cap C, A \cup D)$ and $(B \cap D, A \cup C)$. Analogously to $L(A, C)$ we define the separators $L(A, D)$, $L(B, C)$ and $L(B, D)$.

Observation 4.3. $|L(A, C)| + |L(B, D)| = |A \cap B| + |C \cap D|$. □

Given two separations (A, B) and (C, D) of order at most k such that $L(A, C)$ contains at most k vertices, then the corner separation $(A \cap C, B \cup D)$ has order at most k. If additionally P is a tangle of order $k + 1$ such that A and C are big in P, then the side $A \cap C$ of $(A \cap C, B \cup D)$ is big in P; this follows from the first property of tangles as B, D and $A \cap C$ cover all edges. We shall refer to that property of tangles as the *corner property*.

Another property of tangles of order $k + 1$ that also follows from the property that no three small sides cover is that they are *robust*[8]; that is:

[8]In the context of tangles and separations the term 'robust' is used by different authors

given two separations (A,B) and (C,D), where the first separation has order at most k and the second separation has arbitrary finite order such that the corner separators $L(A,C)$ and $L(B,C)$ have at most $k-1$ vertices. Then if the side C is big, then one of the corners $C\cap A$ or $C\cap B$ must be big.

A separation (A,B) *distinguishes* two tangles P_1 and P_2 if the side A is big in some P_i and small in P_{i+1}. Note that (A,B) distinguishes the P_i if and only if (B,A) distinguishes them. A separation distinguishes P_1 and P_2 *efficiently* if it distinguishes them and has minimal order amongst all separations distinguishing them.

4.2 Blocks and torsos

Given a set $\mathcal{N}$ of separations, an *$\mathcal{N}$-block* is a maximal set of vertices no two of which are separated[9] by a separation in $\mathcal{N}$. For any $\mathcal{N}$-block β, any separation in $\mathcal{N}$ has (at least) one side that includes β. And β can be written as an intersection of all these sides.

Let $\mathcal{N}$ be a nested[10] set of separations of order at most k and let β be an $\mathcal{N}$-block of at least $k+1$ vertices. Let P be a tangle of order $\ell+1$ greater than k. We say that the tangle P *lives* in the block β if for every separation (A,B) of G of order at most k with $\beta\subseteq A$ the side A is big in P.

Remark 4.4. Unlike for finite graphs, not every tangle of order $k+1$ of an infinite graphs lives in an $\mathcal{N}$-block; indeed for tangles defined from ends the intersections of all big sides of separations in $\mathcal{N}$ may be empty. Example 4.5 shows that the definition of 'lives in' cannot be weakened by replacing '(A,B) of G' by '(A,B) of $\mathcal{N}$'.

Example 4.5. In this example we construct a nested set of separations of order three such that the intersections of the big sides of the tangle forms a block of size four in which the tangle does not live. We obtain the graph G from a ray by attaching vertices v and w complete to the ray and then attaching an edge complete to v and w, see Figure 7. The tangle we focus on is the tangle of the vertex-end of G. The set $\mathcal{N}$ consists of those separations of the form $(P_n+v+w, G\setminus P_n)$, where P_n is the initial subpath of the ray of length n. The attached edge together with v and w is an $\mathcal{N}$-block. This

to mean different things that do not seem to be closely related. The notion we use was first defined in [10].

[9] Two vertices v_1 and v_2 are *separated* by a separation (A,B) if some v_i is in $A\setminus B$ and v_{i+1} is in $B\setminus A$.

[10] A *nested set* is a set of separations that are pairwise nested. A separation is *nested* with a set if it is nested with every separation in that set.

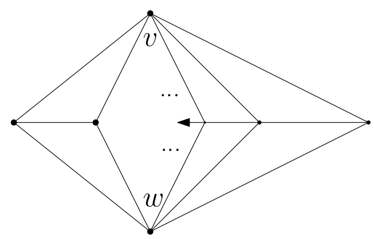

Figure 7: The graph G.

$\mathcal{N}$-block is the intersection the big sides of separations in $\mathcal{N}$. Still the tangle does not live in that block in the sense that it induces a tangle in that block in the sense of Lemma 4.14 below.

The next lemma gives a criterion when tangles do live in blocks.

Lemma 4.6. *Let $\mathcal{N}$ be a nested set of separations of order at most k and (A, B) a separation of order $\ell \geq k+1$ nested with $\mathcal{N}$. Assume that (A, B) distinguishes two tangles P and Q efficiently. Then there is an $\mathcal{N}$-block β such that $A \cap B \subseteq \beta$ and P and Q live in β.*

Proof. Since the separation (A, B) is nested with any separation in $\mathcal{N}$, no such separation separates the vertex set $A \cap B$. Note that $A \cap B$ contains at least $k+1$ vertices. Let β be the unique $\mathcal{N}$-block including $A \cap B$: as above β is unique as each separation in $\mathcal{N}$ has precisely one side containing $A \cap B$.

Next we show that P and Q live in β. For that let (C, D) be a separation of order at most k of G with $\beta \subseteq C$. Our aim is to show that C is big in P and Q.

Suppose for a contradiction that the side D is big in one of the tangles, say P. By symmetry, we may assume that the side A is big in P. Since the link $(A \cap B) \setminus C$ is empty, the corner separation $(A \cap D, B \cup C)$ has order at most k. As P has the corner property, the corner $A \cap D$ is big in P. On the other hand, the side $B \cup C$ must be big in Q as it includes the big side B. Hence the corner separation $(A \cap D, B \cup C)$ distinguishes P and Q. As this separation has order at most k, this is a contradiction to the efficiency of the separation (A, B). Thus the side C is big in both tangles P and Q. $\square$

Observation 4.7. *In the proof of Lemma 4.6 we do not make use of the whole strengths of the property of tangles that three small sides do not cover but just of the slightly weaker corner property. This will be used only once, namely in Observation 4.15*

Given a set $\mathcal{N}$ of separations and an $\mathcal{N}$-block β, the *torso* $G_T[\beta]$ of β is obtained from $G[\beta]$ by adding an edge between any two vertices of β that are in a common separator $A \cap B$ of some separation (A, B) in $\mathcal{N}$. This definition is compatible with the usual definition of torso [14] in the context of tree-decompositions: if $\mathcal{N}$ is the set of separations corresponding to the edges of a tree-decomposition, then the vertex set of every maximal part is an $\mathcal{N}$-block and its torso is just the torso of that part. Moreover, we have the following.

Lemma 4.8. *Let K be a component of $G - \beta$. Then any two vertices v and w in the neighbourhood of K in β are adjacent in the torso $G_T[\beta]$.*

Proof. Let P be a path between v and w whose interior vertices are in K. For each separation (C, D) in $\mathcal{N}$ its restriction to $P \cup \beta$ is $(C \cap (P \cup \beta), D \cap (P \cup \beta))$. By reversing separations in $\mathcal{N}$ if necessary, we may assume that $\beta \subseteq C'$ for every restriction (C', D'). Since nestedness is preserved by restricting, there is one such restriction (C', D') such that $D' \setminus C'$ includes all sets $D'' \setminus C''$ for all other such restrictions (C'', D''). As no vertex of $P - v - w$ is in β the set $D' \setminus C'$ must be equal to $P - v - w$. Hence vw is an edge in the torso. □

Lemma 4.9. *Let $\mathcal{N}$ be a nested set of separations of order at most k and let β be an $\mathcal{N}$-block. Then every component C of $G - \beta$ has at most k neighbours in β.*

Proof. This lemma is a well-known fact for finite graphs[11]. We give an argument that reduces the infinite version to the finite version.

Suppose for a contradiction that some component C of $G - \beta$ has at least $k + 1$ vertices in its neighbourhood. Then there is a finite connected subset C' of C that has a set β' of $k + 1$ vertices of β included in its neighbourhood. We obtain the graph G' from G by deleting all vertices not in the finite vertex set $C' \cup \beta'$. The restrictions $(A', B') = (A \cap G', B \cap G')$ of separations (A, B) in $\mathcal{N}$ form a nested set of separations in G'. Hence we get the desired contradiction by the finite version of the lemma. This completes the proof. □

Remark 4.10. Tangles have many nice properties. However, they do not always induce tangles in blocks they live in, see Example 4.11 below. This

[11] Indeed, let (A, B) be a separation in $\mathcal{N}$ with a vertex v of C contained in the side of (A, B) that does not include β such that the side containing v is inclusion-wise maximal. It is routine to check that the separator $A \cap B$ includes the neighbourhood of the component C.

property will be essential for our proof strategy later on. We will overcome that problem by working within the class of 'robust profiles', a slight superclass of tangles.

It should be noted that robust profiles unlike tangles do not always have the following property, which makes tangles work very well with graph minors: let G' be a minor of G and T' be a tangle in G', then there is a tangle in G inducing T'.[12] This last statement is not used in this paper.

Example 4.11. Consider the unique tangle of order $k+1$ on the complete bipartite graph $K_{k,k+1}$ for $k > 3$. The separations of order k are nested and the torso of the block in which the tangle lives is isomorphic to K_k. However, there is no tangle of order $k+1$ at K_k. (The largest tangle has order roughly $\frac{2}{3} \cdot k$.)

Robust profiles[13] will be defined like tangles except that we weaken the property that three small sides never cover; namely we just forbid this for very particular configurations. To be precise, we define *robust profiles* like 'tangles' except that we replace the first property that three small sides never cover all edges by the following three properties.

1. no two small sides cover all edges;
2. the corner property;
3. the robustness property.

Example 4.12. We have seen above that tangles are examples of robust profiles. A different example is the robust profile of order $k+1$ on the graph K_k.

All definitions for tangles are extended to robust profiles in the obvious way. The proof of Lemma 4.14 is the only one in the paper where we make use of the difference between tangles and robust profiles (except from those implicit places where we apply Lemma 4.14). This is necessary in order to cope with examples such as those in Example 4.11.

Next we define how a robust profile living in an $\mathcal{N}$-block β defines a robust profile in the torso graph $G_T[\beta]$. The *restriction* of a separation (A, B) of G to β is the separation $(A \cap \beta, B \cap \beta)$ of $G[\beta]$.

[12]Conversely, it can be shown that any profile in a graph G' with the property that it induces a profile in any graph G that has G' as a minor is a tangle.

[13]Profiles were introduced in [10]. In that paper 'robust' is called '∞-robust'. The results and proofs of this paper extend verbatim to 'r-robust profiles' for any natural number r. The reader interested in such generalisation is refered to [5], an earlier version of this paper.

Lemma 4.13. *Given an $\mathcal{N}$-block β and a separation (A, B) nested with $\mathcal{N}$, the restriction of (A, B) is a separation in the torso graph $G_T[\beta]$.*

Proof. It suffices to show that for any separation $(C, D) \in \mathcal{N}$ that $C \cap D$ is a subset of A or B. This follows from the nestedness of (A, B) with (C, D). $\square$

For any separation (A', B') of a torso graph $G_T[\beta]$, there is a separation (A, B) of G that restricts to (A', B') and has the same separator. Now let P be a robust profile of order $\ell + 1 > k$ that lives in an $\mathcal{N}$-block β. The *induced robust profile* P_β of P at β is defined as follows. A side A' of a separation (A', B') of the torso graph $G_T[\beta]$ of order at most ℓ is *big in* P_β if and only if there is a side A of a separation (A, B) of G that restricts to (A', B') and has the same separator such that A is big in P.

Lemma 4.14. *Assume that a robust profile P of order $\ell + 1 > k$ lives in the $\mathcal{N}$-block β. Then the induced robust profile P_β is a robust profile of the torso $G_T[\beta]$.*

Proof. First we show that if (A', B') is a separation of order at most k of the torso, then it can have at most one big side in P_β.

By the component property, there is a component K of the graph $G - A' \cap B'$ such that the side $K \cup (A' \cap B')$ is big in P. As P lives in β by assumption, the block β is included in that side. As β has at least $k + 1$ vertices, the component K contains a vertex of the block β. That is, the vertex set $K' = K \cap \beta$ is not empty.

As K' is a restriction of a connected set, it is connected in the torso by Lemma 4.8. As the vertex set K' is disjoint from the separator $A' \cap B'$, there is a unique side of the separation (A', B') that includes K', say A'. Now let (A, B) be any separation of G that restricts to (A', B') and has the separator $A' \cap B'$. Since K includes K', the set K cannot be included in B. So it is included in A. So A must be big in P as it includes a big side. In particular B is small in P. Since (A, B) is arbitrary, B' cannot be big.

To see that P_β has the component property, let X be a set of at most ℓ vertices of the torso. Let K be a component of $G - X$ such that the side $(K \cup X)$ of $(K \cup X, G \setminus K)$ is big in P. Let $K' = K \cap \beta$, which is connected in the torso by Lemma 4.8.

Next we show that K' is not empty. Suppose not for a contradiction. Then the component K contains no vertex of the block β. So K is a component of $G - \beta$. By Lemma 4.9, the component K has at most k neighbours in the block β. So the separation $(K \cup N(K), G \setminus K)$ has order at most k. As P lives in β, the side $G \setminus K$ is big in P. This is a contradiction to the fact that P is a robust profile as the side $K \cup N(K)$ is big in P. Hence the set K' must be nonempty.

Let K'' be the component of the torso $G_T[\beta]$ without X including the connected nonempty set K'. Then $K'' \cup X$ is big in the induced robust profile. So P_β has the component property.

It remains to show that two small sides in the torso do not cover and to show the corner property and robustness for P_β. To see the first, suppose for a contradiction that the torso is covered by two small sides. We observe that if the edge set of a complete graph is covered by two subgraphs, then one of these subgraphs must include the whole vertex set of the complete graph.[14] Hence by Lemma 4.8 for each component K of G without the torso, there is one of the sides that includes the whole neighbourhood of K. So we can assign each component K to a side that includes its neighbourhood. Each of the two covering sides together with its assigned components forms a side of a separation of order at most ℓ in the graph G. As this side restricts to a small side in the torso, it must be small in the original robust profile P by definition of the induced robust profile and by the first part of the proof. Hence the graph G is covered by two sides that are small in P, which is not possible as P is a robust profile P.

Having shown that two small sides cannot cover in the torso, it remains to verify the the corner property and robustness for P_β. In a nutshell, they are both true as taking the corner separation commutes with taking the torso. In detail, let (A', B') and (C', D') be two separations of the torso of order at most ℓ and assume that their corner separator $L(A', C')$ contains at most ℓ vertices. Then there are separations (A, B) and (C, D) of G that restrict to (A', B') and (C', D') and have the same separator; in particular, all vertices of the separators $A \cap B$ and $C \cap D$ are vertices of the torso. Hence the corner separator $L(A, C)$ is equal to the corner separator $L(A', C')$. So the corner property for P_β follows from the corner property for P. Similarly robustness for P_β follows from robustness for P. So P_β is a robust profile of the torso. □

Observation 4.15. *Lemma 4.6 is true with 'tangle' replaced by 'robust profile'.*

Proof. This follows from Observation 4.7. □

[14]This observation is no longer true if we consider covers by three subgraphs instead; and is the reason why this proof does not work for tangles (which is not surprising in view of Example 4.11).

4.3 Extending separations of the torsos

The aim of this subsection is to explain how for a given nested set $\mathcal{N}$ of separations and a torso of an $\mathcal{N}$-block, a nested set of separations of the torso can be extended to a nested set of separations of the whole graph that is nested with $\mathcal{N}$. This is more technical and hence more complicated as one might expect. Indeed, extending a single separation of the torso is quite easy – but it is not uniquely defined. We have to make some choices. If we make these choices arbitrarily for two nested separations, it could happen that their extensions are no longer nested, see Example 4.23 below.

Throughout this subsection we fix a nested set $\mathcal{N}$ of separations and an $\mathcal{N}$-block β. For each separation $(C, D) \in \mathcal{N}$ at least one of the sides C and D includes β. Let $\mathcal{N}_\beta$ consist of those separations (C, D) such that β is included in C and (C, D) or (D, C) is in $\mathcal{N}$.

Given a separation (A, B) of the torso $G_T[\beta]$, one way to 'extend' (A, B) to a separation of G is to decide for each component of $G - \beta$, whether we put it on the A-side or on the B-side. Below we define what it means when such a component is 'forced'. Informally, it is forced when we must put it on the A-side in order to extend (A, B) to a separation of G.

A component K of $G-\beta$ is *forced at step* 1 by (A, B) if one of its vertices has a neighbour in $A \setminus B$. A separation $(C, D) \in \mathcal{N}_\beta$ is *forced at step* $2n+2$ if there is a component K forced at step $2n+1$ that contains a vertex of $D \setminus C$. A component K of $G - \beta$ is *forced at step* $2n+1$ for $n > 0$ if there is a separation $(C, D) \in \mathcal{N}_\beta$ forced at step $2n$ so that K contains a vertex of $D \setminus C$. An alternative definition of 'forcing' is the following.

Example 4.16. We define the bipartite graph whose left side are the components of $G - \beta$ and whose right side are the separations in $\mathcal{N}_\beta$. We add an edge between a component K and a separation (C, D) if K contains a vertex of $D \setminus C$. A component (or separation) is forced if and only if its connected component in this bipartite graph contains a component forced at step one. We will not use the fact that this definition is equivalent in our proofs.

The following lemma implies that if a component is forced at some step, it is forced at step one or three; and if a separation is forced, it is forced at step two or four.

Lemma 4.17. *Let (C, D) be a separation in $\mathcal{N}_\beta$ forced by (A, B). There is some (C', D') in $\mathcal{N}_\beta$ forced by (A, B) with $D \setminus C \subseteq D' \setminus C'$ such that some vertex of $C' \cap D'$ is in $A \setminus B$.*

Proof. Let $2n$ be the smallest step at which (C, D) is forced. We prove Lemma 4.17 by induction on n.

The base case is that $2n = 2$. Let K be a component forced at step one 'forcing' (C, D); here we say that K *forces* the separation (C, D) if there is a vertex of K in $D \setminus C$ and (C, D) is not forced at an earlier step than K.

As K is forced at step one, there is a vertex v of K that has a neighbour w in $A \setminus B$. As v is not in β, there is a separation (E, F) in $\mathcal{N}_\beta$ such that v is in $F \setminus E$. As w is in β, it is in E. As it has a neighbour in $F \setminus E$, it also must be in F.

We call a separation (E', F') a *candidate* if the separator $E' \cap F'$ contains a vertex of $A \setminus B$ and $F' \setminus E'$ contains a vertex of the component K. For example, the separation (E, F) is a candidate. To conclude the base case, we show the following.

Sublemma 4.18. *Assume that there is a candidate. Then there is a separation (C', D') in $\mathcal{N}_\beta$ forced by (A, B) with $D \setminus C \subseteq D' \setminus C'$ such that some vertex of $C' \cap D'$ is in $A \setminus B$.*

Proof. We pick a candidate (E, F). Let w be a vertex of the separator $E \cap F$ in $A \setminus B$. If the vertex w was in the separator $C \cap D$, the lemma would be true with '(C, D)' in place of '(C', D')'. Hence we may assume that the vertex w of β is not in the side D as $\beta \subseteq C$. So the vertex w is in the link $(E \cap F) \setminus D$. So from the nestedness of (C, D) and (E, F) it follows that either $D \setminus C \subseteq F \setminus E$ or else $D \setminus C$ and $F \setminus E$ are vertex-disjoint.

Our aim is to construct a candidate (E, F) that satisfies the first condition $D \setminus C \subseteq F \setminus E$. Let u and v be vertices of K that are in $D \setminus C$ and $F \setminus E$, respectively (such vertices exist as we may assume that $D \setminus C$ is not empty and (E, F) is a candidate).. Let P be a path from the vertex u to vertex v included in the component K of $G - \beta$. By assumption for every vertex x on P, there is a separation $(C_x, D_x) \in \mathcal{N}_\beta$ with $x \in D_x \setminus C_x$. If possible we choose the separation (C_x, D_x) such that the vertex w is in the separator (in that case it is a candidate).

Our goal is to show that it is possible to choose the separation (C_u, D_u) at u such that the vertex w is in the separator. Indeed, then we can use the nestedness of (C_u, D_u) and (C, D) to deduce as above that $D \setminus C \subseteq D_u \setminus C_u$ or else $D \setminus C$ and $D_u \setminus C_u$. However, here the second outcome is not possible as the vertex u is in the intersection of these two sets.

Suppose for a contradiction that such a choice for (C_u, D_u) is not possible. Let x be the vertex on the path P nearest to u such that the vertex w is in its separator $C_x \cap D_x$. Let y be the neighbour of x on P nearer to

u, which exists as $x \neq u$. Then the vertex w is in the link $(C_x \cap D_x) \setminus D_y$. As above we deduce from the nestedness of the separations (C_x, D_x) and (C_y, D_y), that either $D_y \setminus C_y \subseteq D_x \setminus C_x$ or else $D_y \setminus C_y$ and $D_x \setminus C_x$ are vertex-disjoint. Since we cannot choose (C_x, D_x) in place of (C_y, D_y), the first outcome is impossible. The second outcome is not possible either as the vertex $x \in D_x \setminus C_x$ is adjacent to the vertex $y \in D_y \setminus C_y$. So this is the desired contradiction. Hence we can choose (C_u, D_u) such that it is a candidate, which completes the proof as shown above. $\square$

So the base case follows from Sublemma 4.18 and the fact that (E, F) is a candidate.

Now let $n > 1$ and assume that we already proved Lemma 4.17 for separations (E, F) forced at some step before $2n$. Let (C, D) be a separation in $\mathcal{N}_\beta$ forced at step $2n$. Let K force (C, D). Let (C_1, D_1) be a separation forcing K, which exists as $n > 1$. By the induction hypothesis, there is a separation (C_2, D_2) in $\mathcal{N}_\beta$ forced by (A, B) with $D_1 \setminus C_1 \subseteq D_2 \setminus C_2$ such that some vertex w of $C_2 \cap D_2$ is in $A \setminus B$. As $D_1 \setminus C_1 \subseteq D_2 \setminus D_2$, there is a vertex v of K that is in $D_2 \setminus C_2$. So (C_2, D_2) is a candidate. So the induction step follows from Sublemma 4.18. This completes the proof. $\square$

Lemma 4.19. *For any separation (A, B) of the torso, no component K of $G - \beta$ is forced by both (A, B) and (B, A).*

Proof. As any component of $G - \beta$ forces some separation in $\mathcal{N}_\beta$, it suffices to show that no separation (C, D) in $\mathcal{N}_\beta$ is forced by both (A, B) and (B, A). Suppose for a contradiction that there is such a separation (C, D).

By Lemma 4.17, there is a separation $(C', D') \in \mathcal{N}_\beta$ forced by (A, B) with $D \setminus C \subseteq D' \setminus C'$ such that some vertex v of $C' \cap D'$ is in $A \setminus B$. As $D' \setminus C'$ is a superset of $D \setminus C$, the separation (C', D') is also forced by (B, A). Applying Lemma 4.17 to (C', D') and to (B, A), yields a separation (C'', D'') with $D' \setminus C' \subseteq D'' \setminus C''$ such that some vertex w of $C'' \cap D''$ is in $B \setminus A$. Since the vertex v is in β, it must be in C''. As D'' includes D', it also is in D''. In short, v is in the separator $C'' \cap D''$.

Hence the separation (C'', D'') witnesses that vw is an edge of the torso. As v is in $A \setminus B$ and w is in $B \setminus A$, we deduce that (A, B) cannot be a separation of the torso. That is the desired contradiction. $\square$

Having finished the proof of Lemma 4.19, we now define naive extensions of separations of the torso, explain why they are not quite the object we need and define extensions of nested sets of separations of the torso.

Given a separation (A, B) of the torso, the side $\hat{A}$ is obtained from A by adding all components K of $G - \beta$ that are forced at some step. We obtain $\hat{\hat{B}}$ from B by adding all components K that are not forced at any step. Note that $\hat{\hat{B}} = B \cup (G \setminus \hat{A})$. We define the *naive extension* of (A, B), denoted by $\widehat{(A, B)}$, to be $(\hat{A}, \hat{\hat{B}})$. This construction ensures that $\widehat{(A, B)}$ is a separation of G that restricts to (A, B).

Remark 4.20. We chose the notation '$\hat{\hat{B}}$' instead of simply '$\hat{B}$' as the term $\hat{A}$ for the separation (A, B) and the term '$\hat{\hat{A}}$' for the separation (B, A) need not a priori agree – and in fact they do not agree for (A, B) defined as in Example 4.23.

In particular, the reverse separation of $\widehat{(A, B)}$ is in general not equal to $\widehat{(B, A)}$.

Observation 4.21. *Given two separations (A, B) and (X, Y) of the torso, if $(A, B) \leq (X, Y)$, then $\widehat{(A, B)} \leq \widehat{(X, Y)}$.* □

Observation 4.22. *Let (A, B) be a separation of the torso and $(C, D) \in \mathcal{N}_\beta$ be a proper separation. Then $\widehat{(A, B)}$ or its reverse separation $(\hat{\hat{B}}, \hat{A})$ is $\leq (C, D)$.*

In particular, $\widehat{(A, B)}$ is nested with every proper separation of $\mathcal{N}$.

Proof. First assume that the separation (C, D) is forced at some step. Then $D \setminus C$ is a subset of $\hat{A}$. Since the separator $A \cap B$ of the separation $(\hat{\hat{B}}, \hat{A})$ is a subset of the block β, which is included in the side C, we conclude that $D \setminus C$ is a subset of $\hat{A} \setminus \hat{\hat{B}}$. As in the proof of Observation 2.5 one combines this with the assumption that the separation (C, D) is proper to deduce that $(D, C) \leq \widehat{(A, B)}$.

Hence it remains to consider the case that the separation (C, D) is not forced at any step. Analoguously as above, one shows that $(D, C) \leq (\hat{\hat{B}}, \hat{A})$ in that case. So $\widehat{(A, B)}$ or its reverse separation $(\hat{\hat{B}}, \hat{A})$ is $\leq (C, D)$.

□

Example 4.23 gives an example of nested separations (A, B) and (C, D) of the torso $G_T[\beta]$ whose naive extensions $\widehat{(A, B)}$ and $\widehat{(C, D)}$ are not nested.

Example 4.23. Let G be the labelled graph depicted in Figure 8. The set $\mathcal{N}$ consists of the separation of order one and its reverse. Then the

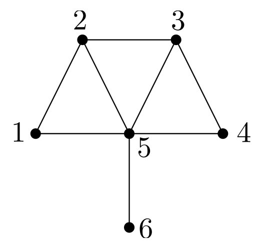

Figure 8: The graph G is obtained by from a triangle by attach two more triangles at distinct edges and by then attaching a leaf at the unique vertex of degree four.

torso is $G-6$. We define $A=\{1,2,3,5\}$, $B=\{3,4,5\}$, $C=\{2,3,4,5\}$, $D=\{1,2,5\}$. Then (A,B) and (C,D) are nested but not $\widehat{(A,B)}$ and $\widehat{(C,D)}$.

Examples like Example 4.23 motivate the slightly technical definition of $\widetilde{\mathcal{L}}$ below. Given a nested set $\mathcal{L}$ of separations of $G_T[\beta]$, the *extension* $\widetilde{\mathcal{L}}$ of $\mathcal{L}$ (depending on a well-order $((A_\alpha,B_\alpha)\mid\alpha\in\kappa)$ of $\mathcal{L}$) is the set $\{\widetilde{(A_\alpha,B_\alpha)}\mid(A_\alpha,B_\alpha)\in\mathcal{L}\}$, where the *extension* $\widetilde{(A,B)}$ of (A,B) is defined as follows: for the smallest element (A_0,B_0) of the well-order, we just let $\widetilde{(A_0,B_0)}=\widehat{(A_0,B_0)}$.

Assume that we already defined $\widetilde{(A_\alpha,B_\alpha)}$ for all $\alpha<\gamma$. A component K of $G-\beta$ is γ*-forced* if there is some $\alpha<\gamma$ such that K is a subset of $\tilde{B}_\alpha$ and $(B_\alpha,A_\alpha)\le(A_\gamma,B_\gamma)$. We obtain $\tilde{A}_\gamma$ from A_γ by adding all components K of $G-\beta$ that are forced by (A_γ,B_γ) or are γ-forced. We obtain $\tilde{B}_\gamma$ from B_γ by adding all other components. The extension $\widetilde{(A_\gamma,B_\gamma)}$ of (A_γ,B_γ) is defined to be $(\tilde{A}_\gamma,\tilde{B}_\gamma)$.

Example 4.24. The nested set $\mathcal{L}=\{(A,B),(C,D)\}$ defined in Example 4.23 has different extensions $\widetilde{\mathcal{L}}$ depending on which well-order we choose.

Observation 4.25. *The extension $(\tilde{A}_\gamma,\tilde{B}_\gamma)$ is a separation.*

Proof. It suffices to show that any component K of $G-\beta$ included in $\tilde{A}_\gamma$ is not forced by (B_γ,A_γ). By Lemma 4.19, we may assume that K is γ-forced. As any class of ordinals has a least element, there is some α minimal such that K is a subset of $\tilde{B}_\alpha$ and $(B_\alpha,A_\alpha)\le(A_\gamma,B_\gamma)$. In particular, K is not forced by (A_α,B_α). As $B_\gamma\setminus A_\gamma$ is a subset of $A_\alpha\setminus B_\alpha$, we deduce that K cannot be forced by (B_γ,A_γ). $\square$

Observation 4.26. *Any separation (A, B) of the torso has the same separator as its extension $\widetilde{(A, B)}$.* $\square$

Observation 4.27. *For any two separations (A, B) and (B, A) in $\mathcal{L}$, the extension $\widetilde{(A, B)}$ is the reverse of $\widetilde{(B, A)}$.*

Proof. We may assume that $(A, B) = (A_\alpha, B_\alpha)$ and $(B, A) = (A_\gamma, B_\gamma)$ for some $\alpha < \gamma$. It suffices to show that $\tilde{B}_\alpha = \tilde{A}_\gamma$. By construction $\tilde{B}_\alpha \subseteq \tilde{A}_\gamma$. Suppose for a contradiction that $\tilde{B}_\alpha$ is a proper subset of $\tilde{A}_\gamma$. Then there is a component K that is included in $\tilde{A}_\alpha$ and in $\tilde{A}_\gamma$. We split into four cases and derive a contradiction in each of them.

Case 1A: K is forced by (A, B) and (B, A). This is impossible by Lemma 4.19.

Case 1B: K is forced by (A, B) and γ-forced. So there is some ordinal $\delta < \gamma$ such that K is a subset of $\tilde{B}_\delta$ and $(B_\delta, A_\delta) \leq (A_\gamma, B_\gamma)$. Then $(A, B) \leq (A_\delta, B_\delta)$. So K is forced by (A_δ, B_δ). So it cannot be a subset of $\tilde{B}_\delta$, a contradiction.

Case 2A: K is α-forced and forced by (B, A). This case is analogue to Case 1B.

Case 2B: K is α-forced and γ-forced. So there is some ordinal $\delta < \alpha$ such that K is a subset of $\tilde{B}_\delta$ and $(B_\delta, A_\delta) \leq (A_\alpha, B_\alpha)$; and there is some ordinal $\epsilon < \gamma$ such that K is a subset of $\tilde{B}_\epsilon$ and $(B_\epsilon, A_\epsilon) \leq (A_\gamma, B_\gamma)$. To summarise:

$$(B_\delta, A_\delta) \leq (A, B) \leq (A_\epsilon, B_\epsilon)$$

If $\delta < \epsilon$, then the component K is ϵ-forced and hence not in $\tilde{B}_\epsilon$, which is impossible. Similarly, we also cannot have $\delta > \epsilon$. So $\delta = \epsilon$. But then the component K is included in the sides B_δ and $A_\delta \supseteq B_\epsilon$, which is the desired contradiction. $\square$

Observation 4.28. *For any two separations (A_α, B_α) and (A_γ, B_γ) in the nested set $\mathcal{L}$, their extensions $\widetilde{(A_\alpha, B_\alpha)}$ and $\widetilde{(A_\gamma, B_\gamma)}$ are nested.*

Proof. By symmetry we may assume that $\alpha < \gamma$. If B_α is a subset of A_γ, it follows immediately from the definitions that $\widetilde{(B_\alpha, A_\alpha)} \leq \widetilde{(A_\gamma, B_\gamma)}$. If A_α is a subset of A_γ, then $\widetilde{(A_\alpha, B_\alpha)} \leq \widetilde{(A_\gamma, B_\gamma)}$ by construction.

The other two cases can be deduced using Observation 4.27 as follows. First assume that (B_γ, A_γ) is not in $\mathcal{L}$. Then we add that separation at the end of the well-order for $\mathcal{L}$. Now we apply the above argument to (A_α, B_α) and (B_γ, A_γ). Hence $\widetilde{(A_\alpha, B_\alpha)}$ and $\widetilde{(B_\gamma, A_\gamma)}$ are nested. By Observation 4.27 also $\widetilde{(A_\alpha, B_\alpha)}$ and $\widetilde{(A_\gamma, B_\gamma)}$ are nested.

The same argument works if (B_γ, A_γ) is in $\mathcal{L}$ but in the well-order after position α. If it is before α, we replace (A_α, B_α) or (A_γ, B_γ) by their reverses if they appear before in the well-order and then do the above argument. This implies the desired result by Observation 4.27. $\square$

Observation 4.29. *For any separation $(A, B) \in \mathcal{L}$, its extension $\widetilde{(A, B)}$ is nested with every proper separation in $\mathcal{N}$.*

Proof. Let $(A, B) = (A_\gamma, B_\gamma)$. We say that a separation (C, D) of $\mathcal{N}_\beta$ is γ*-forced* if there is some component K of $G - \beta$ that contains a vertex of $D \setminus C$ and is γ-forced or forced by (A_γ, B_γ).

We claim that if a separation of $\mathcal{N}_\beta$ is γ-forced, then every component of $G - \beta$ that contains a vertex of $D \setminus C$ is γ-forced or forced by (A_γ, B_γ). Indeed, if any such component is forced by a separation of the nested set $\mathcal{L}$, then all of these components are. Hence this follows by transfinite induction on the well-order of $\mathcal{L}$.

Using this, we can argue as in the proof of Observation 4.22. $\square$

Lemma 4.30. *Let $\mathcal{N}$ be a nested set of proper separations and let β and γ be distinct $\mathcal{N}$-block. Let $\mathcal{L}_\beta$ and $\mathcal{L}_\gamma$ be nested sets of separations of $G_T[\beta]$ and $G_T[\gamma]$, respectively. Then $\widetilde{\mathcal{L}}_\beta$ is a set of nested separations. For any separations $(A, B) \in \mathcal{L}_\beta$ and $(C, D) \in \mathcal{L}_\gamma$, their extensions $\widetilde{(A, B)}$ and $\widetilde{(C, D)}$ are nested. Moreover, they are nested with every separation in $\mathcal{N}$.*

Proof. The set $\widetilde{\mathcal{L}}_\beta$ is nested by Observation 4.28. The 'Moreover'-part follows from Observation 4.29. So it remains to show that for any separations $(A, B) \in \mathcal{L}_\beta$ and $(C, D) \in \mathcal{L}_\gamma$, the extensions $\widetilde{(A, B)}$ and $\widetilde{(C, D)}$ are nested.

Since the blocks β and γ are distinct, there is a separation (E, F) of $\mathcal{N}$ such that one side includes the block β and the other side includes the block γ. By symmetry we may assume that β is included in E and γ is included in F. By Observation 4.29 $\widetilde{(A, B)}$ is nested with (E, F). An argument as in the proof of Observation 4.29 gives that $F \setminus E$ is included in $B \setminus A$ or $A \setminus B$. So either $\widetilde{(A, B)}$ or its reverse is $\leq (E, F)$.

By Observation 4.27 it would be enough to show that one of $\widetilde{(A, B)}$ or its reverse is nested with $\widetilde{(C, D)}$. Hence by replacing '(A, B)' by (B, A) if necessary, we assume that $\widetilde{(A, B)} \leq (E, F)$. Similarly, one may assume that $(E, F) \leq \widetilde{(C, D)}$. Combining this yields that $\widetilde{(A, B)}$ and $\widetilde{(C, D)}$ are nested. $\square$

Observation 4.31. *Let β, P, Q, P_β and Q_β as in Observation 4.15. Let $\mathcal{L}$ be a nested set of separations in $G_T[\beta]$. If a separation $(C,D) \in \mathcal{L}$ distinguishes the induced robust profiles P_β and Q_β in the torso $G_T[\beta]$, then the extension $\widetilde{(C,D)}$ distinguishes the robust profiles P and Q.*

Proof. By symmetry we may assume that C is big in P_β and D is big in Q_β. As P_β is a robust profile by Lemma 4.14, the component property yields that there is a component K_1 of $\beta - (C \cap D)$ such that $K_1 \cup (C \cap D)$ is big in P_β. So K_1 is a subset of C.

The extension $\widetilde{(C,D)} = (\tilde{C}, \tilde{D})$ has the separator $C \cap D$ by Observation 4.26. Let K_1' be the components of $G - C \cap D$ such that the side $K_1' \cup (C \cap D)$ is big in P. As P_β is induced by P, it must be that K_1 contains a vertex of K_1'. In particular K_1' cannot be a subset of $\tilde{D}$. So it must be a subset of $\tilde{C}$. Hence $\tilde{C}$ is big in P. Similarly one shows that $\tilde{D}$ is big in Q. So the extension $\widetilde{(C,D)}$ distinguishes the robust profiles P and Q. $\square$

4.4 Miscellaneous

The lemmas summarised in this subsection are well-known.

Lemma 4.32. *Let (A,B) and (C,D) be proper separations such that $A \setminus B$ is connected and does not intersect the separator $C \cap D$. Then (A,B) and (C,D) are nested.*

Proof. By the definition of nestedness, it suffices to show that $(A,B) \leq (C,D)$ or $(A,B) \leq (D,C)$. As the connected set $A \setminus B$ does not intersect the separator $C \cap D$, it is included in $C \setminus D$ or $D \setminus C$. By symmetry, we may assume that is is included in $C \setminus D$. So $A \setminus B$ is included in $C \setminus D$. Hence by Observation 2.5 (A,B) and (C,D) are nested. $\square$

Lemma 4.33 ([9, Lemma 2.2])**.** [15] *Let (A,B), (C,D) and (E,F) be proper separations such that first (A,B) and (C,D) are not nested and second the corner separation $(A \cap C, B \cup D)$ is not nested with (E,F). Then (E,F) is not nested with (A,B) or (C,D).*

A separation (A,B) of a graph G is *tight* if every component of G without the separator $A \cap B$ has the whole separator $A \cap B$ in its neighbourhood.

Lemma 4.34. *Let (A,B) be a separation of order at most k. Let (C,D) be a tight separation such that the graph G without the separator $C \cap D$ has at*

[15] This is Lemma 2.2 of that paper with the roles of '(C,D)' and '(E,F)' interchanged.

least $k+1$ components. Then one of the links $(C \cap D) \setminus A$ or $(C \cap D) \setminus B$ is empty.

Proof. Suppose not for a contradiction, then there are vertices $v \in (C \cap D) \setminus A$ and $w \in (C \cap D) \setminus B$. Then v and w are in the neighbourhood of every component of G without the separator $C \cap D$. Thus there are $k+1$ internally disjoint paths from v to w. All of these paths contain vertices of the separator $A \cap B$. This contradicts the assumption that the separator $A \cap B$ contains at most k vertices. $\square$

Given two vertices v and w, a separator S separates v and w *minimally* if each component of $G - S$ containing v or w has the whole of S in its neighbourhood.

Lemma 4.35 ([24, Statement 2.4])**.** *Given vertices v and w and $k \in \mathbb{N}$, there are only finitely many distinct separators of size at most k separating v from w minimally.*

5 Distinguishing the tangles

The aim in this section is to construct for any graph a nested set of separations of finite order that distinguishes any two vertex-ends efficiently, which is needed in the proof of Theorem 1. A related result is proved in [11]. Actually, we shall prove the stronger statement that there is a nested set $\mathcal{N}$ of separations that distinguishes any two tangles efficiently. A simplified version of this proof for finite graphs has been published in [7].

Overview of the proof

We shall construct the set $\mathcal{N}$ as an ascending union of sets $\mathcal{N}_k$ one for each $k \in \mathbb{N}$, where $\mathcal{N}_k$ is a nested set of separations of order at most k distinguishing efficiently any two tangles[16] of order $k+1$, see Figure 9. Any two tangles of order $k+2$ that are not distinguished by $\mathcal{N}_k$ will live in the same $\mathcal{N}_k$-block. We obtain $\mathcal{N}_{k+1}$ from $\mathcal{N}_k$ by adding for each $\mathcal{N}_k$-block β a nested set $\widetilde{\mathcal{N}_{k+1}}(\beta)$ that distinguishes efficiently any two tangles of order $k+2$ living in β. Working in the torsos $G_T[\beta]$ will ensure that the sets $\widetilde{\mathcal{N}_{k+1}}(\beta)$ for different blocks β will be nested with each other.

Summing up, we are left with the task of finding in these torso graphs $G_T[\beta]$ a nested set distinguishing efficiently tangles of order $k+2$. Theorem 5.2 deals with this problem if the torso $G_T[\beta]$ is 'nice enough'. In order

[16] Actually this is not quite correct as we need 'robust profiles' instead of 'tangles'. This detail will be discussed at the end of the sketch.

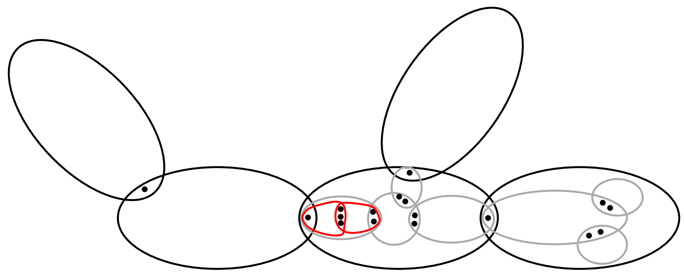

Figure 9: The tree-decomposition corresponding to the nested set $\mathcal{N}_1$ is indicated by black parts. In torsos of that tree-decomposition, in grey we indicated a tree-decomposition for $\mathcal{N}_2$. In each torso of that tree-decomposition we have a further tree-decomposition given by $\mathcal{N}_3$, etc.

to make all torso graphs nice enough, we first do an additional step in which we enlarge $\mathcal{N}_k$ a little bit so that for the larger nested set the new torso graphs are the old ones with the junk cut off. The main lemma for this enlargement is Lemma 5.3.

As explained in Remark 4.10 and Example 4.11[17], for such a torso-approach to work we need to work within the superclass of robust profiles that includes all the tangles (instead of just the tangles).

Finishing the overview, we first state Theorem 5.2 and Lemma 5.3 and introduce the necessary definitions for that.

For any robust profile P and $k \in \mathbb{N}$, its *restriction* P_k to k consists of those separations in P that have order at most k. The order of P_k is the minimum of $k+1$ and the order of P. A *(robust) profile set* is a set of robust profiles that is closed under restrictions. Until the end of Subsection 5.2, we fix a graph G, a number $k \in \mathbb{N} \cup \{\infty\}$ and a profile set $\mathcal{P}$.

A nested set $\mathcal{N}$ of separations is *extendable* (for $\mathcal{P}$) if for any two (distinct) robust profiles in $\mathcal{P}$ of the same order, there is some separation distinguishing these two robust profiles efficiently that is nested with $\mathcal{N}$.

A separation is *relevant* (for a number $k \in \mathbb{N} \cup \{\infty\}$, a graph G and a profile set $\mathcal{P}$) if it has order at most k and it distinguishes two robust profiles in $\mathcal{P}$ efficiently – in particular, it has finite order. We denote the set of all relevant separations by $R(k, \mathcal{P}, G)$.

Given a separation (A, B), a component C of $G - A \cap B$ is *degenerated* if its neighbourhood[18] $N(C)$ is a proper subset of the separator $A \cap B$, see

[17]Indeed, the robust profile induced by any tangle in the torso need not be a tangle.

[18]Throughout, we denote the neighbourhood of a vertex set C by $N(C)$.

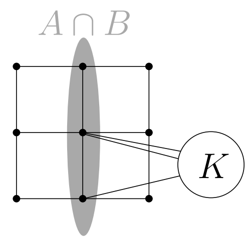

Figure 10: The separator $A \cap B$ is indicated in grey and has size three. The component K of $G - A \cap B$ has only two neighbours in the separator $A \cap B$ and thus is degenerated.

Figure 10.

A separation is *degenerated* relative to (A, B) if it is of the form $(C \cup N(C), G \setminus C)$, where C is a degenerated component of $G - A \cap B$. Given a set $\mathcal{S}$ of separations, its *degenerator* is the set of separations that are degenerated relative to some separation in $\mathcal{S}$. We denote the degenerator of the set $R(k, \mathcal{P}, G)$ of relevant separations by $S(k, \mathcal{P}, G)$. If it is clear from the context what G is, we shall just write $R(k, \mathcal{P})$ or $S(k, \mathcal{P})$, or even just $R(k)$ or $S(k)$.

Example 5.1. Every relevant separation in $R(k)$ is tight if and only if $S(k)$ is empty.

Theorem 5.2. *Let $k \in \mathbb{N}$. Assume that $S(k) = \emptyset$ and $R(k-1) = \emptyset$. Let $\mathcal{N}$ be any nested subset of $R(k)$ that is inclusion-wise maximal.*

Then $\mathcal{N}$ distinguishes any two robust profiles of order $k+1$ in $\mathcal{P}$ efficiently and is extendable.

Lemma 5.3. *If $R(k-1)$ is empty, then the degenerator $S(k)$ is a nested extendable set of separations.*

5.1 Proof of Lemma 5.3.

In this subsection we prove Lemma 5.3. First we need some preparation.

A separation (A, B) *pre-disqualifies* a separation (C, D) if the order of (C, D) is strictly larger than the sizes $|L(A, C)|$ and $|L(B, C)|$ of corner separators. A separation (A, B) *disqualifies* a separation (C, D) if it pre-disqualifies (C, D) or its reverse (D, C).

The following lemma shows that relevant separations cannot be disqualified.

Lemma 5.4. *If (C, D) distinguishes robust profiles P_1 and P_2 efficiently, then no separation (A, B) disqualifies (C, D).*

Proof. We may assume that C is big in P_1 and D is big in P_2. Suppose for a contradiction that some separation (A, B) pre-disqualifies (C, D).

So the order of (C, D) is strictly larger than $|L(A, C)|$ and $|L(B, C)|$. The side $B \cup D$ of the corner separation $(A \cap C, B \cup D)$ is big in the robust profile P_2 as it includes a big side. By the efficiency of (C, D), this corner separation cannot distinguish P_1 and P_2. Thus $A \cap C$ is small in P_1. A similar argument shows that also the corner $B \cap C$ must be small in P_1. This violates the robustness of P_1. This is a contradiction to the assumption that P_1 is a robust profile. Hence (A, B) cannot pre-disqualify (C, D). Analogously, one shows that (A, B) cannot pre-disqualify (D, C). □

Lemma 5.5. *Let (A, B) and (C, D) be two separations distinguishing robust profiles in $\mathcal{P}$ efficiently such that the order of (A, B) is k and the order of (C, D) is at least k. Let K be a degenerated component of $G - A \cap B$.*

If $R(k-1)$ is empty, then K does not intersect the separator $C \cap D$.

Proof. By symmetry, we may assume that the component K is included in $A \setminus B$.

Sublemma 5.6. *If $R(k-1)$ is empty, then the side $K \cup N(K)$ is small in every robust profile of order greater than k of G.*

Proof. By assumption, there is a robust profile P of order greater than k such that the side B is big in P. As the side $G \setminus K$ of the separation $(K \cup N(K), G \setminus K)$ includes the big side B, it must also be big in P. Since $R(k-1)$ is empty, the side $K \cup N(K)$ is small in every robust profile of order greater than k. □

Sublemma 5.7. *If the corner separation $(A \cap C, B \cup D)$ distinguishes two robust profiles of order greater than k efficiently, then the component K does not intersect the separator $C \cap D$.*

Proof. The corner separation of the separations $(A \cap C, B \cup D)$ and $(G \setminus K, N(K) \cup K)$ is $((A \cap C) \setminus K, B \cup D \cup K)$. In particular, $((A \cap C) \setminus K, B \cup D \cup K)$ is a separation. The separator of the separation $((A \cap C) \setminus K, B \cup D \cup K)$ is the corner separator $L(A, C)$ without K; in formulas $L(A, C) \setminus K$. This separation also distinguishes two robust profiles efficiently: if the side $B \cup D$ is big in a robust profile, then also the side $B \cup D \cup K$ is big in that robust profile. By Sublemma 5.6 if the side $A \cap C$ is big in a robust profile, then also the side $(A\cap C)\setminus K$ is big in that robust profile as it satisfies the component property and $G\setminus K$ is big in all robust profiles. Hence by the efficiency of the separation $(A\cap C, B\cup D)$, it must be that the order of the separation $((A\cap C)\setminus K, B\cup D\cup K)$ is at least the order of $(A\cap C, B\cup D)$; that is, the corner separator $L(A,C)$ does not intersect the component K.

If the component K intersects the separator $C\cap D$, it does so in the link $(C\cap D)\setminus B$ as K is a subset of $A\setminus B$. Since this link is a subset of the corner separator $L(A,C)$, the component K cannot intersect that link as shown above. So the component K does not intersect the separator $C\cap D$. $\square$

Let Q_1 and Q_2 be two robust profiles distinguished efficiently by the separation (C,D) such that C is big in Q_1 and D is big in Q_2. By replacing the separation (A,B) by the separation $(B\cup K, A\setminus K)$ if necessary we may assume that the side A is big in Q_1.

Sublemma 5.8. *Either $|L(A,C)|\leq |C\cap D|$ and the corner $A\cap C$ is big in Q_1 or else $|L(A,D)|\leq |C\cap D|$ and the corner $A\cap D$ is big in Q_2.*

Proof. Either the side A or B must be big in the robust profile Q_2. We distinguish two cases.

Case 1: the side B is big in Q_2.

If $|L(B,D)| < |A\cap B|$, then the corner $B\cap D$ is in Q_2 by the corner property. Thus the corner separation $(B\cap D, A\cup C)$ will distinguish Q_1 and Q_2, which is impossible by the efficiency of (C,D). Thus by Observation 4.3 $|L(A,C)|\leq |C\cap D|$, yielding that the corner $A\cap C$ is big in Q_1 by the corner property, as desired.

Case 2: the side A is big in Q_2.

By Lemma 5.4, the separation (C,D) does not pre-disqualify (B,A). Thus either $|L(B,C)|\geq |A\cap B|$ or $|L(B,D)|\geq |A\cap B|$. In the first case, by Observation 4.3 $|L(A,D)|\leq |C\cap D|$. Then the corner $A\cap D$ is big in Q_2 by the corner property. Similarly in the second case, $|L(A,C)|\leq |C\cap D|$. Then the corner $A\cap C$ is big in Q_1 by the corner property, as desired. $\square$

By Sublemma 5.8, one of the corner separations $(A\cap C, B\cup D)$ or $(A\cap D, B\cup C)$ distinguishes the robust profiles Q_1 and Q_2 efficiently. Hence by Sublemma 5.7 or the corresponding fact for the corner separation $(A\cap D, B\cup C)$, we deduce that the component K does not intersect the separator $C\cap D$.

$\square$

Proof of Lemma 5.3. Let (A,B) be a relevant separation in $R(k)$ and (C,D) be some separation that distinguishes two robust profiles efficiently of order at least k. Let K_1 be a degenerated component of $G - A \cap B$ and K_2 be a component of $G - C \cap D$. In order to see that $S(k)$ is a nested, it suffices to show that for any such K_1 and K_2 that the separations $(K_1 \cup N(K_1), G \setminus K_1)$ and $(K_2 \cup N(K_2), G \setminus K_2)$ are nested. This is true by Lemma 5.5 and Lemma 4.32. In order to see that $S(k)$ is an extendable, it suffices to show that for any such K_1 and (C, D) that the separations $(K_1 \cup N(K_1), G \setminus K_1)$ and (C, D) are nested. This is true by Lemma 5.5 and Lemma 4.32, as well. $\square$

5.2 Proof of Theorem 5.2.

We actually prove the following extension of Theorem 5.2. It is more general in the sense that it allows for even more flexibility which sets $\mathcal{N}$ we could choose. Recall that a separation is relevant (in $R(\infty)$) if it distinguishes some two robust profiles efficiently.

Theorem 5.9. *Let $k \in \mathbb{N}$. Assume that $S(k) = \emptyset$ and $R(k-1) = \emptyset$. Any set $\mathcal{N}$ of nested tight separations of order at most k that are not disqualified by any relevant separation is extendable.*

In particular, any maximal such set distinguishes any two robust profiles of order $k+1$ in $\mathcal{P}$ efficiently.

Before we prove Theorem 5.9, we need some intermediate lemmas. Throughout this subsection, we assume that $S(k)$ is empty. Let U be the set of those tight separations of order at most k that are not disqualified by any relevant separation. Since $R(k)$ is a subset of U, Theorem 5.9 implies Theorem 5.2.

Lemma 5.10. *For any relevant separation (A, B) such that $A \setminus B$ is connected, there are only finitely many separation $(C, D) \in U$ not nested with (A, B).*

Proof. First, we show that the separation (A, B) is nested with every separation $(C, D) \in U$ such that the link $(A \cap B) \setminus C$ is empty. By Lemma 4.32, it suffices to show that the link $(C \cap D) \setminus B$ is empty. As (A, B) does not pre-disqualify (D, C), one of the links $(C \cap D) \setminus A$ or $(C \cap D) \setminus B$ is empty. As we are done otherwise, we may assume that the link $(C \cap D) \setminus A$ is empty. If (C, D) is not nested with (A, B), there must be a component of K of $G - (C \cap D)$ all of whose neighbours are in the center $(A \cap B) \cap (C \cap D)$. As (C, D) is tight, it must be that $(C \cap D) = (A \cap B) \cap (C \cap D)$ so that $(C \cap D) \setminus B$ is empty. Hence (A, B) and (C, D) are nested by Lemma 4.32.

Similarly one shows that the separation (A, B) is nested with every separation $(C, D) \in U$ such that the link $(A \cap B) \setminus D$ is empty.

It remains to show that there are only finitely many separations $(C, D) \in U$ not nested with (A, B). As shown above, in that case both links $(A \cap B) \setminus C$ and $(A \cap B) \setminus D$ are nonempty. By Lemma 4.35, there are only finitely many triples (v, w, T) where $v, w \in (A \cap B)$ and T is a separator of size at most k separating v and w minimally. Since each separator $C \cap D$ for some (C, D) as above is such a separator T, it suffices to show that there are only finitely many separations in U that have the same separator as (C, D). This is true as the connected[19] graph G without the separator $C \cap D$ has only finitely many components by Lemma 4.34 (in fact it has at most $|C \cap D| + 1$ components). $\square$

Lemma 5.11. *Let $\mathcal{N}$ be a nested subset of U. For any two robust profiles P and Q of order $\ell \geq k+1$ that are not distinguished by any separation of order less than k, there is some separation (A, B) that is nested with $\mathcal{N}$ and distinguishes P and Q efficiently.*

Proof. First, we show that there is a separation (A, B) distinguishing P and Q efficiently that is nested with all but finitely many separations of $\mathcal{N}$. Since $S(k)$ is empty, $R(k)$ is a subset of U. Let (A, B) be a separation distinguishing P and Q efficiently. As the robust profiles P and Q have the component property, we can pick (and we do pick) the separation (A, B) such that $A \setminus B$ is connected. By Lemma 5.10, (A, B) is nested with all but finitely many separations of $\mathcal{N}$. Hence we can pick a separation (A, B) distinguishing P and Q efficiently such that it is not nested with a minimal number of $(C, D) \in \mathcal{N}$.

Suppose for a contradiction that there is some separation $(C, D) \in \mathcal{N}$ that is not nested with (A, B). We may assume that (C, D) does not distinguish P and Q since otherwise (C, D) would distinguish P and Q efficiently by assumption. Thus either the side C is big in both P and Q or else the side D is big in both P and Q. Since (D, C) is nested with $\mathcal{N}$, we may by symmetry assume that C is big in both P and Q.

Since (A, B) does not pre-disqualify (D, C) by the definition of U, either $|L(A, D)| \geq |C \cap D|$ or $|L(B, D)| \geq |C \cap D|$. By symmetry, we may assume that $|L(A, D)| \geq |C \cap D|$. By exchanging the roles of P and Q if necessary, we may assume that A is big in P and B is big in Q. By Observation 4.3, $|L(B, C)| \leq |A \cap B|$. Note that the corner $B \cap C$ is small in P as it is included in the small side B of P. On the other hand, by the corner property the corner $B \cap C$ is be big in Q. Thus the corner separation $(B \cap C, A \cup D)$

[19]Recall that the assumption that $S(k-1)$ is empty implies that the graph G is connected.

distinguishes P and Q efficiently. Any separation in $\mathcal{N}$ not nested with the corner separation $(B \cap C, A \cup D)$ is by Lemma 4.33 not nested with (A, B). As (C, D) is nested with the corner separation $(B \cap C, A \cup D)$, this corner separation violates the minimality of (A, B). Hence (A, B) is nested with $\mathcal{N}$, completing the proof. □

Proof of Theorem 5.9. By Lemma 5.11 and since $R(k-1)$ is empty, any nested subset $\mathcal{N}$ of U is extendable.

Since by assumption any relevant separation in $R(k)$ is in the set U, it follows that any maximal such set $\mathcal{N}$ distinguishes any two robust profiles of order $k+1$ in $\mathcal{P}$. It distinguishes efficiently as $R(k-1)$ is empty. □

5.3 Proof of the main result of this section.

In this subsection, we prove the following.

Theorem 5.12. *For any graph G, there is a nested set $\mathcal{N}$ of separations that distinguishes efficiently any two robust profiles (that are not restrictions of one another).*

First we need an intermediate lemma about sticking together a nested set $\mathcal{N}$ of proper separations with nested sets of separations in the torsos of the $\mathcal{N}$-blocks. We fix a finite number k and a profile set $\mathcal{P}$. Let $\mathcal{N}$ be a nested set of separations of order at most k that is extendable for $\mathcal{P}$ and that distinguishes efficiently any two robust profiles of $\mathcal{P}$ that can be distinguished by a separation of order at most k in G. For each $\mathcal{N}$-block β, we denote by $\mathcal{P}(\beta)$ the set of robust profiles in $\mathcal{P}$ living in β. And let $\mathcal{N}_\beta$ be a set of nested separations of the torso $G_T[\beta]$ of β that is extendable for the induced robust profiles, induced by those robust profiles in $\mathcal{P}(\beta)$. We abbreviate $\mathcal{M} = \mathcal{N} \cup \bigcup \widetilde{\mathcal{N}_\beta}$, where the union ranges over all $\mathcal{N}$-blocks β. (Here in order to define the sets $\widetilde{\mathcal{N}_\beta}$ we choose arbitrary well-orderings on the sets $\mathcal{N}_\beta$.)

Lemma 5.13. *The set $\mathcal{M}$ is nested, proper and extendable for $\mathcal{P}$.*

Proof. The set $\mathcal{M}$ is nested by Lemma 4.30. The separations in $\mathcal{N}$ are proper by assumption and those in some $\mathcal{N}_\beta$ are proper as they are efficient. By Observation 4.26 extensions of proper separations are proper.

It remains to show for every $\ell \geq k+1$ and any two robust profiles P and Q in $\mathcal{P}$ that are distinguished efficiently by a separation of order ℓ in G that there is a separation nested with $\mathcal{M}$ that distinguishes P and Q efficiently. We may assume that P and Q both have order $\ell+1$ as $\mathcal{P}$ is a profile set.

Since $\mathcal{N}$ is extendable, there is a separation (A, B) of order ℓ nested with $\mathcal{N}$ that distinguishes P and Q. By Observation 4.15, there is a unique $\mathcal{N}$-block β including the separator $A \cap B$ such that P and Q live in β. The restriction of (A, B) to β distinguishes the robust profiles P_β and Q_β, induced by P and Q respectively. As $\mathcal{N}_\beta$ is extendable, there is a separation (A', B') of the torso $G_T[\beta]$ that distinguishes P_β and Q_β efficiently; in particular it has order at most ℓ. By Lemma 4.30, the extension $\widetilde{(A', B')}$ is nested with $\mathcal{M}$. By Observation 4.26 it has order at most ℓ. So by Observation 4.31 it distinguishes P and Q efficiently. As P and Q were arbitrary, the nested set $\mathcal{M}$ is extendable. □

Proof of Theorem 5.12. We shall construct the nested set $\mathcal{N}$ of Theorem 5.12 as a nested union of sets $\mathcal{N}_k$ one for each $k \in \mathbb{N} \cup \{-1\}$, where $\mathcal{N}_k$ is a nested extendable set of separations of order at most k that distinguishes any two robust profiles efficiently that are distinguished by a separation of order at most k. We start the construction with $\mathcal{N}_{-1} = \emptyset$. Assume that we already constructed $\mathcal{N}_k$ with the above properties.

We denote by $\mathcal{P}$ the set of all robust profiles of G. For an $\mathcal{N}_k$-block β, we denote the set of robust profiles in $\mathcal{P}$ living in β by $\mathcal{P}(\beta)$, and by $\mathcal{P}_\beta$ the induced robust profiles of β, induced by robust profiles in $\mathcal{P}(\beta)$. Note that $\mathcal{P}_\beta$ is a profile set by Lemma 4.14.

Sublemma 5.14. *The set $R(k, \mathcal{P}_\beta, G_T[\beta])$ is empty.*

Proof. Suppose for a contradiction, two robust profiles P_β and Q_β in $\mathcal{P}_\beta$ can be distinguished by a separation (A, B) of order at most k. Then $\widetilde{(A, B)}$ has the same order as (A, B) by Observation 4.26. It distinguishes the robust profiles P and Q which induce P_β and Q_β by Observation 4.31. Since P_β and Q_β are distinct, also P and Q are distinct. But then by the induction hypothesis P and Q are distinguished by $\mathcal{N}_k$ – as they are distinguishable by the separation $\widetilde{(A, B)}$ of order at most k. This contradicts the fact that P and Q are both in $\mathcal{P}(\beta)$. □

By Sublemma 5.14, we can apply Lemma 5.3 to the torso graph $G_T[\beta]$ and $\mathcal{P}_\beta$, yielding that the degenerator $S(k + 1, \mathcal{P}_\beta, G_T[\beta])$ is a nested extendable set of separations. For each $S(k+1, \mathcal{P}_\beta, G_T[\beta])$-block β', we define $\mathcal{P}(\beta')$ and $\mathcal{P}_{\beta'}$ similarly as $\mathcal{P}(\beta)$ and $\mathcal{P}_\beta$, respectively.

Sublemma 5.15. *The degenerator $S(k + 1, \mathcal{P}'_\beta, G_T[\beta'])$ is empty.*

Proof. Suppose for a contradiction that this degenerator is not empty. Then there is a relevant separation (C, D) in $R(k+1, \mathcal{P}'_\beta, G_T[\beta'])$ that has a degenerated component. By Observation 4.26 and Observation 4.31, the extension $\widetilde{(C, D)}$ of (C, D) distinguishes efficiently two robust profiles in $\mathcal{P}(\beta)$. In particular that extension is relevant in $G_T[\beta]$, that is, it is contained in $R(k+1, \mathcal{P}_\beta, G_T[\beta])$.

By assumption there is a degenerated component K' of β' without the separator $C \cap D$. The component K' is included in a component K of β without the separator $C \cap D$. As K' and K have the same neighbours in the separator $C \cap D$, also K is degenerated. So the separation $(K \cup N(K), \beta \setminus K)$ is in the degenerator $S(k+1, \mathcal{P}_\beta, G_T[\beta])$. Thus β' is disjoint from the component K. So K' is empty, which is the desired contradiction. $\square$

By Zorn's Lemma we pick a maximal nested subset $\mathcal{N}(\beta')$ of $R(k+1, \mathcal{P}_{\beta'}, G_T[\beta'])$, that is, of separations of order at most $k+1$ in the graph $G_T[\beta']$ distinguishing efficiently two robust profiles in $\mathcal{P}_{\beta'}$. By Theorem 5.2 the set $\mathcal{N}(\beta')$ is extendable for $\mathcal{P}_{\beta'}$ and distinguishes any two robust profiles of order $k+2$ in $\mathcal{P}_{\beta'}$ efficiently.

Let $\mathcal{N}_{k+1}(\beta)$ be the union of the nested set $S(k+1, \mathcal{P}_\beta, G_T[\beta])$ with the sets $\widetilde{\mathcal{N}(\beta')}$, where β' is an $S(k+1, \mathcal{P}_\beta, G_T[\beta])$-block. By Lemma 5.13, $\mathcal{N}_{k+1}(\beta)$ is a nested and extendable set of separations of order at most $k+1$ in $G_T[\beta]$. Let $\mathcal{N}_{k+1}$ be the union of $\mathcal{N}_k$ with the sets $\widetilde{\mathcal{N}_{k+1}(\beta)}$, where β is an $\mathcal{N}_k$-block. By applying Lemma 5.13 again, we deduce that $\mathcal{N}_{k+1}$ is a nested and extendable set of separations of order at most $k+1$ in G.

Sublemma 5.16. *$\mathcal{N}_{k+1}$ distinguishes efficiently any two robust profiles P and Q of G that are distinguished by a separation of order at most $k+1$.*

Proof. As $\mathcal{N}_k$ is a subset of $\mathcal{N}_{k+1}$, we may assume by the induction hypothesis that any separation distinguishing P and Q efficiently has order $k+1$. Let (C, D) be such a separation distinguishing them efficiently. As $\mathcal{N}_{k+1}$ is extendable by Lemma 5.13, we can pick (and we do pick) (C, D) so that it is nested with $\mathcal{N}_{k+1}$.

By Observation 4.15, there is an $\mathcal{N}_k$-block β including the separator $C \cap D$ such that P and Q live in β. The induced robust profiles of P and Q in β are denoted by P_β and Q_β, respectively. The restriction of (C, D) is nested with the degenerator $S(k+1, \mathcal{P}_\beta, G_T[\beta])$ by construction. By Observation 4.15, there is an $S(k+1, \mathcal{P}_\beta, G_T[\beta])$-block β' including the separator $C \cap D$ such that P and Q induce distinct robust profiles in $\mathcal{P}_{\beta'}$. These induced robust profiles are distinguished efficiently by $\mathcal{N}(\beta')$ by construction.

Applying Observation 4.31 twice and Observation 4.26 yields that P and Q are distinguished by $\mathcal{N}_{k+1}$. As every separation in $\mathcal{N}_{k+1}$ has order at most $k+1$, the robust profiles P and Q are distinguished efficiently by $\mathcal{N}_{k+1}$. $\square$

Finally, the nested union $\mathcal{N}$ of the sets $\mathcal{N}_k$ is a nested set of separations that distinguishes efficiently any two robust profiles of the same order, as desired. $\square$

Corollary 5.17. *For any graph G, there is a nested set $\mathcal{N}$ of finite separations that contains for any two vertex-ends a separation distinguishing them efficiently.*

Proof. By Example 4.2, each vertex-end induces a tangle, which in return defines a robust profile. All these tangles are distinct (also as robust profiles) for different vertex-ends. So this is a consequence of Theorem 5.12. $\square$

6 A tree-decomposition distinguishing the topological ends

In this section, we prove Theorem 1 already mentioned in the introduction. A key lemma in the proof of Theorem 1 is the following.

Lemma 6.1. *Let G be a graph with a finite nonempty set W of vertices. Then G has a star-decomposition*[20] *$(S, Q_s|s \in V(S))$ of finite adhesion such that each topological end lives in a part Q_s with s a leaf.*

Moreover, only the central part Q_c contains vertices of W, and for each leaf s, a topological end lives in the part Q_s, and the set $Q_s \setminus Q_c$ is connected.

Proof that Lemma 6.1 implies Theorem 1. We shall recursively construct a sequence $\mathcal{T}^n = (T^n, P_t^n|t \in V(T^n))$ of tree-decompositions of G of finite adhesion as follows. We start by picking a vertex r' of G arbitrarily and we obtain $\mathcal{T}^1$ by applying Lemma 6.1 with $W = \{r'\}$. We refer to r' as the *rooting vertex*. Assume that we already constructed $\mathcal{T}^n$. For each leaf s of $\mathcal{T}^n$, we denote by W_s the set of those vertices in Q_s also contained in some other part of $\mathcal{T}^n$. Note that W_s is contained in the part adjacent to Q_s and thus is finite. By Lemma 6.1, we obtain a star-decomposition $\mathcal{T}_s$ of $G[Q_s]$ such that no $w \in W_s$ is contained in a leaf part of $\mathcal{T}_s$ and such that each topological end living in Q_s lives in a leaf of $\mathcal{T}_s$. We obtain $\mathcal{T}^{n+1}$ from $\mathcal{T}^n$ by replacing each leaf part Q_s by $\mathcal{T}_s$, which is well-defined as the set W_s is contained in a unique part of $\mathcal{T}_s$.

[20] A *star-decomposition is a tree-decomposition, where the decomposition tree is a star.*

By r, we denote the center of $\mathcal{T}_1$. For each $j < m < n$, the balls of radius j around r in T^m and T^n are the same. Thus we take T to be the tree whose nodes are those that are eventually a node of T^n. For each node $t \in V(T)$, the parts P_t^n are the same for n larger than the distance between t and r, and we take P_t to be the limit of the P_t^n.

It is easily proved by induction that each vertex in the set W_s for s a leaf of the tree T^n has distance at least $n-1$ from the rooting vertex r' in the graph G. Thus for each $j < n$ the ball of radius j around the rooting vertex r' in G is included in the union over all parts P_t^n where t is in the ball of radius j around r in $\mathcal{T}_n$. Hence $(T, P_t|t \in V(T))$ is a tree-decomposition, and it has finite adhesion by construction.

It remains to show that the ends of T define precisely the topological ends of G, which is done in the following four sublemmas.

Sublemma 6.2. *Each topological end ω of G lives in an end of T.*

Proof. There is a unique leaf s of T^n such that ω lives in P_s^n. Let s_n be the predecessor of s in T^n. Then ω lives in the end of T to which $s_1 s_2 \dots$ belongs. □

Sublemma 6.3. *In each end τ of T, there lives a vertex-end of G.*

Proof. Let X be a spanning tree of the graph G. Our aim is to find a ray included in X whose vertex-end lives in the end τ.

Let $s_1 s_2$... be the ray in T starting at r that belongs to the end τ. By construction, the sets W_{s_i} are disjoint and finite. Let U be the union of the sets W_{s_i}. Since each vertex is separated by some set W_{s_i} from all but finitely many vertices of U, the tree X does not include a subdivision of an infinite star with all leaves in U. Hence by the Star-Comb-Lemma[21] [14, Section 8], there is a comb with infinitely many leaves in the set U. Thus the vertex end of the ray of that comb lives in the end τ. □

Sublemma 6.4. *No two distinct vertex-ends ω_1 and ω_2 of G live in the same end τ of T.*

Proof. Suppose for a contradiction, there are such vertex-ends ω_1, ω_2 living in the same end τ. Let U be a finite separator separating ω_1 from ω_2 and let n be the maximum over the distances between the rooting vertex r' and a vertex in U. Let s be the unique node of the tree T on the ray starting

[21] The Star-Comb-Lemma says that if U is an infinite vertex set in a tree X, then either X contains a subdivision of an infinite star with all leaves in U or X contains a comb with all leaves in U; here a comb is obtained from a ray by attaching a path at each vertex.

at the root belonging to the end τ that has distance n from the root. By construction, in the tree T^{n+1} the node s is leaf. Let C_i be the component of the graph $G - U$ in which the vertex-end ω_i lives. Recall that the leaf-part Q_s (of $\mathcal{T}_{n+1}$) with the separator W_s removed is connected. Since the set W_s separates the separator U from the set $Q_s \setminus W_s$, the connected set $Q_s \setminus W_s$ is contained in a component of the graph $G - U$. As the vertex-end ω_i lives in the graph $Q_s \setminus W_s$ by assumption, it must be that the set $Q_s \setminus W_s$ is a subset of the component C_i. Hence the components C_1 and C_2 intersect, which is the desired contradiction. □

Sublemma 6.5. *No vertex u dominates a vertex-end ω living in some end of T.*

Proof. Suppose for a contradiction a vertex u dominates the vertex-end ω. Let n be the distance between the vertex u and the rooting vertex r' in G. Then there is a leaf s of the tree T^{n+1} such that the vertex-end ω lives in the leaf-part Q_s. Thus the finite set W_s separates the vertex u from the vertex-end ω, contradicting the assumption that the vertex u dominates the vertex-end ω. □

Sublemma 6.2, Sublemma 6.3, Sublemma 6.4 and Sublemma 6.5 imply that the ends of T define precisely the topological ends of G, as desired. □

Remark 6.6. Let $(T, \leq)$ be the tree order on T as in the proof of Theorem 1 where the root r is the smallest element. We remark that we constructed $(T, \leq)$ such that $(T, P_t | t \in V(T))$ has the following additional property. For each edge tu with $t \leq u$, the vertex set $\bigcup_{w \geq u} V(P_w) \setminus V(P_t)$ is connected.

Moreover, we construct $(T, P_t | t \in V(T))$ such that if st and tu are edges of T with $s \leq t \leq u$, then $V(P_s) \cap V(P_t)$ and $V(P_t) \cap V(P_u)$ are disjoint.

In order to prove Lemma 6.1, we need the following.

Lemma 6.7. *Let G be a connected graph and W a finite and nonempty vertex set of G. Then there is a set $\mathcal{X}$ of separations (A_i, B_i) of finite order such that every vertex-end not dominated by a vertex of W lives in a side B_i. Moreover, the sets $B_i \setminus A_i$ are disjoint and the set W is vertex-disjoint from all sides B_i.*

Proof that Lemma 6.7 implies Lemma 6.1. If a set $B_i \setminus A_i$ has several components, we replace the separation (A_i, B_i) in $\mathcal{X}$ by the set of separations for the form $(G \setminus C, C \cup N(C))$, where C is a component of $B_i \setminus A_i$. Call the resulting set $\mathcal{X}'$. By replacing $\mathcal{X}$ by $\mathcal{X}'$ if necessary, we may assume that all sets $B_i \setminus A_i$ in Lemma 6.7 are connected.

We may assume that the graph G in Lemma 6.1 is connected. Our aim is to construct a star-decomposition of G. It gets a leaf-part for every every separation (A_i, B_i) in $\mathcal{X}$ such that a topological end lives in B_i. Its part is B_i.

Let C be the intersection of the sides A_i with $(A_i, B_i) \in \mathcal{X}$ together with all sides B_i such that no topological end lives in B_i. The set C gets the central part of our star-decomposition. By construction, this is a star-decomposition and it has finite adhesion. This star-decomposition has all the desired properties by construction. □

The rest of this section is devoted to the proof of Lemma 6.7. We shall need the following lemma.

Lemma 6.8. *Let G be a connected graph and W a finite nonempty vertex set. There is a nested set $\mathcal{N}$ of proper separations (A_i, B_i) of finite order such that every vertex-end not dominated by a vertex of W lives in a side B_i and the set W is vertex-disjoint from all sides B_i.*

Moreover, if two distinct separations (A_i, B_i) and (A_j, B_j) of $\mathcal{N}$ satisfy $(A_i, B_i) \leq (A_j, B_j)$, then the order of (A_i, B_i) is strictly larger than the order of (A_j, B_j).

Proof. We obtain the graph G_W from the graph G by first deleting the vertex set W and then adding a copy of the complete graph[22] K_ω in such a way that it is joined completely to the neighbourhood of W in G. Applying Corollary 5.17 to the graph G_W, yields a nested set $\mathcal{N}'$ of separations of finite order such that any two vertex-ends of G_W are distinguished efficiently by a separation in $\mathcal{N}'$. Let τ be the vertex-end to which the rays of the newly added copy of K_ω belong. Let $\mathcal{N}''$ consist of those separations in $\mathcal{N}'$ that distinguish τ efficiently from some other vertex-end. By reversing separations in $\mathcal{N}''$ if necessary, we may assume that the added graph K_ω is included in the side A for every separation $(A, B) \in \mathcal{N}''$. As the separations in $\mathcal{N}''$ distinguish efficiently, for no separation (A, B) in $\mathcal{N}''$ the side B contains a vertex of the added graph K_ω.

Given a natural number k, a *k-sequence* $((A_\alpha, B_\alpha)|\alpha \in \gamma)$ (for $\mathcal{N}''$) is an ordinal indexed sequence of elements of $\mathcal{N}''$ of order at most k such that if $\alpha < \beta$, then $B_\alpha \subseteq B_\beta$. (Recall that every separation in $\mathcal{N}''$ is proper so $B_\alpha \subseteq B_\beta$ implies that $(A_\beta, B_\beta) \leq (A_\alpha, B_\alpha)$ by Observation 2.5.) The *union*

[22]By K_ω we denote the complete graph on countably many vertices.

separation of a k-sequence $((A_\alpha, B_\alpha)|\alpha \in \gamma)$ is the separation obtained by taking the union over all B-sides and the intersection of all A-sides, formally it is: $(\bigcap_{\alpha\in\gamma} A_\alpha, \bigcup_{\alpha\in\gamma} B_\alpha)$.

The set $\mathcal{N}'''$ consists of all union separations of k-sequences of $\mathcal{N}''$ for all k. Since we allow constant sequences, the nested set $\mathcal{N}''$ is included in the set $\mathcal{N}'''$. A standard transfinite induction argument yields that the set $\mathcal{N}'''$ is nested.[23] Given a natural number k, the set $\mathcal{N}_k$ consists of those separations of the nested set $\mathcal{N}'''$ that have order at most k; and the set $\mathcal{N}'_k$ consists of those elements (A, B) of the nested set $\mathcal{N}_k$ whose B-side is inclusion-wise maximal in $\mathcal{N}_k$.

We take $\mathcal{N}_W$ to be the union of the nested sets $\mathcal{N}'_k$. By construction, for each separation (A, B) in $\mathcal{N}_W$, the side B contains no vertex of the added graph K_ω. We obtain $\mathcal{N}$ from $\mathcal{N}_W$ by replacing each separation (A, B) in $\mathcal{N}_W$ by the separation where we modify the side A by replacing the added graph K_ω by the finite vertex set W. Clearly, the set $\mathcal{N}$ is a nested set of proper[24] separations of the graph G.

We claim that the nested set $\mathcal{N}$ has all the properties stated in Lemma 6.8: the 'Moreover'-part is clear by construction. Thus it remains to show that each vertex-end ω of G not dominated by a vertex of W lives in a side B_i for some separation (A_i, B_i) in the nested set $\mathcal{N}$.

Let R be a ray belonging to the vertex-end ω. Since the vertex-end ω is not dominated by any vertex of W, for each vertex x of W there is a finite vertex set S_x separating a subray R_x of R from x. We let S be the union of these finite separators S_x. The finite set S separates the intersection R' of the subrays R_x from vertex set W in the graph G. In the graph $G - W$, the set S separates the ray R' from the added graph K_ω. Let ω' be the vertex-end of G_W to which the ray R' belongs. Note that the separator S witnesses that the vertex-end ω' is not equal to the vertex-end τ of the added K_ω. Thus there is a separation (A_i, B_i) of $\mathcal{N}'''$ so that the vertex-end ω' lives in the side B_i. Let k be the order of the separation (A_i, B_i). By Zorn's lemma, the nested $\mathcal{N}'''$ contains a separation (A', B') with $B_i \subseteq B'$ of order at most k whose B-side is inclusion-wise maximal amongst all separations of $\mathcal{N}'''$ of order at most k. By construction the separation (A', B') is in the nested set

[23] Given two members of $\mathcal{N}'''$, by replacing their underlying sequences by different sequences with the same union, one may assume that all members of their sequences are nested with each other in the same way (of the four ways separations could be nested). Hence by transfinite induction also their unions must be nested in that way.

[24] This properness follows from the fact that (A, B) comes from a separation of G_W distinguishing two tangles efficiently and the modifications made by going from G_W to G preserve being proper.

$\mathcal{N}'_k$ and includes a subray of R'. So the separation $((A' \setminus K_\omega) \cup W, B')$ is in $\mathcal{N}$ and the vertex-end ω lives in the side B'. This completes the proof.[25] $\square$

Next we show how Lemma 6.8 implies Lemma 6.7. A good candidate for the nested set $\mathcal{X}$ of Lemma 6.7 might be the separations (A, B) in the nested set $\mathcal{N}$ such that the side B is inclusion-wise maximal amongst members of the nested set $\mathcal{N}$. However, there might be an infinite strictly increasing – that is, the B-sides are strictly increasing – sequence of members in $\mathcal{N}$, whose orders are also strictly increasing, so that we cannot expect that the union of all these B-sides is a side of a finite order separation, and hence cannot come from a separation in $\mathcal{N}$, see Example 6.9. Thus we have to make a more sophisticated choice for $\mathcal{X}$ than just taking the 'maximal members' of $\mathcal{N}$.

Example 6.9. The set W just consists of a single vertex, which is complete to a ray. At each initial path P_n of the ray, we attach a copy of the ladder of width n. The set $\mathcal{N}$ consists of those separations (A_n, B_n) with separator P_n separating the first n ladders from the vertex of W. These separations are strictly increasing in order and this sequence does not have a maximal element in the sense explained above this example.

Proof that Lemma 6.8 implies Lemma 6.7. Let $\mathcal{N}$ be a nested set as in Lemma 6.8. The first step in the proof is to define a graph H that visualises the structure of the nested set $\mathcal{N}$.

Let (A, B) be a separation of the nested set $\mathcal{N}$ such that there is another separation (C, D) in $\mathcal{N}$ with $B \subseteq D$. (Recall that all separations of the set $\mathcal{N}$ are proper. So this implies that $(C, D) \leq (A, B)$ by Observation 2.5.) Then the order of (C, D) is larger than that of (A, B).

Such a separation (C, D) is called a *successor* of the side B (we use the condensed notation 'of B' instead of 'of (A, B)' as every proper separation is uniquely determined by one of its sides); the separation (C, D) is an *immediate succesor* if it has minimal order amongst all successors. Let H be the digraph with vertex set $\mathcal{N}$ where we put in the directed edge BD if the separation (C, D) is an immediate successor of B. A *connected component of H*, is a connected component of the underlying graph of H. A typical

[25] We sketch an alternative proof. First one shows that we may assume that every separation (A, B) in $\mathcal{N}''$ has the property that $A \setminus B$ is connected. Then instead of referring to $\mathcal{N}'''$ one can use the following fact: let P be a tangle of order $k + 1$ and let (A_i, B_i) be a sequence of separations of order at most k distinguishing tangles efficiently such that each A_i is big in P and the sets $A_i \setminus B_i$ are connected and $A_{i+1} \subseteq A_i$. Then the sequence has only finitely many distinct members.

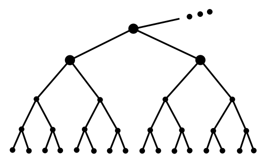

Figure 11: The canopy tree. We obtain an interesting example for the graph H by directing all edges away from the leaves.

example for a connected component of the graph H is the canopy tree, see Figure 11.

Sublemma 6.10. *Let (A', B') and (C', D') be separations in $\mathcal{N}$. Then $B' \subseteq D'$ if and only if there is a directed path in H from B' to D'. Moreover, if two separations (A, B) and (C, D) in $\mathcal{N}$ are not joined by a directed path, then $B \setminus A$ and $D \setminus C$ are disjoint.*

Proof. Clearly, if there is a directed path from B' to D', then $B' \subseteq D'$. Conversely, let (A', B') and (C', D') be separations in $\mathcal{N}$ with $B' \subseteq D'$. Let (A_n, B_n) be a sequence of distinct separations in $\mathcal{N}$ such that $B' \subseteq B_1 \subseteq ... \subseteq B_n \subseteq D'$. By Lemma 6.8, $n \leq |\partial(D')| - |\partial(B')| + 1$. Thus there is a maximal such chain (E_n, F_n), which satisfies $F_1 = B'$ and $F_n = D'$ and $F_{i+1} \in F(B_i)$ for all i between 1 and $n-1$. Hence $F_1 ... F_n$ is a path from B' to D'.

To see that "Moreover"-part, let (A, B) and (C, D) be separations in $\mathcal{N}$. Since the set W is nonempty, the side B does not include the side C. As the sides B and D cannot be subsets of one another by assumption, the nestedness yields that $B \subseteq C$ and $D \subseteq A$. Hence the sets $D \setminus C$ and $B \setminus A$ are disjoint. $\square$

Sublemma 6.11. *Each vertex v of H has out-degree at most one.*[26]

Proof. Suppose for a contradiction the side v has out-degree at least 2. Then there are distinct immediate successors (A, B) and (C, D). By the conditions of Lemma 6.8, it must be that neither $B \subseteq D$ nor $D \subseteq B$. Thus $B \setminus A$ and $D \setminus C$ are disjoint by Sublemma 6.10. Since $v \subseteq B \cap D$, it cannot be the side of a proper separation. This is the desired contradiction to the assumption that v is a side of a separation in $\mathcal{N}$. $\square$

[26] We do not use it in our proof but it follows from this lemma that H is a forest. Indeed, any cycle included in H must by the outdegree condition be a directed cycle. This, however, is impossible by Sublemma 6.10.

Sublemma 6.12. *Any undirected path P joining two vertices v and w contains a vertex u such that vPu and wPu are directed paths which are directed towards u.*

Proof. It suffices to show that the path P contains at most one vertex of out-degree zero on P. If it contained two such vertices, then between them would be a vertex of out-degree two, which is impossible by Sublemma 6.11. $\square$

We define the set $\mathcal{X}$ as the union of sets $\mathcal{X}_K$, one for each component K of H. The sets $\mathcal{X}_K$ are defined as follows. If a component K contains a vertex v_K of out-degree 0, then by Sublemma 6.12 K cannot contain a second such vertex and for any other vertex v in the component K, there is a directed path from v to the vertex v_K directed towards v_K. Hence the side v_K includes any other the side v that is a vertex of the component K. We choose for $\mathcal{X}_K$ the unique separation in $\mathcal{N}$ with the side v_K; here the uniqueness follows from the fact that proper separations are uniquely determined by one of their sides.

Otherwise, every vertex of the component K has outdegree precisely one by Sublemma 6.11. Since the component K cannot contain a directed cycle by Sublemma 6.10, it must contain a directed ray; that is, a ray $B_1B_2\dots$ with $B_i \subsetneq B_{i+1}$. In this case, we define the set $\mathcal{X}_K$ to consist of the separations (C_i, D_i) defined as follows. Let (A_i, B_i) be the unique proper separation with side B_i. We let $(C_1, D_1) = (A_1, B_1)$. Roughly, we obtain (C_i, D_i) from (A_i, B_i) by flipping the set $B_{i-1} \setminus A_{i-1}$ from the side B_i to the side A_i; in formulas for $i > 1$, we let $C_i = A_i \cup (B_{i-1} \setminus A_{i-1})$ and $D_i = B_i \setminus (B_{i-1} \setminus A_{i-1})$. Note that the order of (C_i, D_i) is bounded by the sum of the orders of (A_i, B_i) and (A_{i-1}, B_{i-1}), and thus finite. Since no side B_i contains a vertex of W, the same is true for the sides D_i. This completes the definition of the set $\mathcal{X}_K$ and thus $\mathcal{X}$.

Any two distinct separations (A, B) and (C, D) in the set $\mathcal{X}$ have the property that $B \setminus A$ and $D \setminus C$ are disjoint; indeed, if these separations are in the same set $\mathcal{X}_K$, this is clear by construction. Otherwise it follows from the definition of the side D_i and Sublemma 6.10. Thus it remains to prove the following:

Sublemma 6.13. *Every vertex-end ω not dominated by some vertex of W lives in some side B with $(A, B) \in \mathcal{X}$.*

Proof. By Lemma 6.8, there is a separation (E, F) in the nested set $\mathcal{N}$ such that the vertex-end ω lives in F. Let K be the component of H containing the vertex F. If $\mathcal{X}_K = \{v_K\}$, then $F \subseteq v_K$; and we are done as the vertex-end ω lives in the side v_K. Otherwise let the B_i and the D_i be as in the construction of $\mathcal{X}_K$. If $F = B_j$ for some j, then we pick j minimal such that ω lives in B_j. Since ω does not live in B_{j-1}, it must live in D_j, as desired.

Thus we may assume that the side F is not equal to any side B_j. Let P be a path joining the vertex F and the vertex $B_1 = D_1$. By Sublemma 6.12, the path P contains a vertex u such that the subpaths FPu and B_1Pu are directed paths which are directed towards u. Thus $F \subseteq u$. Since the out-degree is at most one, the path B_1Pu is a subpath of the ray $B_1B_2\dots$. Thus the vertex u is equal to B_j for some j. In particular, $F \subseteq B_j$.

As the side B_j is the union of the finitely many sides $D_1, D_2, ..., D_j$, the vertex-end ω has to live in some side D_i with $i \leq j$. This completes the proof. $\square$

$\square$

Finally we deduce Corollary 2.6.

Proof that Theorem 1 implies Corollary 2.6. By Theorem 1, G has a tree-decomposition $(T, P_t|t \in V(T))$ of finite adhesion such that the ends of T define precisely the topological ends of T, and we choose this tree-decomposition as in Remark 6.6. In particular, we can pick a root r of T such that for each edge tu with $t \leq u$, the vertex set $\bigcup_{w \geq u} V(P_w) \setminus V(P_t)$ is connected.

Furthermore for each edge tu with $t \leq u$, one may assume that if a vertex is in the separator $V(P_t) \cap V(P_u)$, then it has a neighbour in $P_u \setminus P_t$ by deleting other vertices from the part P_u if necessary.

Thus for each such edge tu, there is a finite connected subgraph S_u of the induced subgraph $G[\bigcup_{w \geq u} V(P_w)]$ that contains the separator $V(P_t) \cap V(P_u)$. Let Q_t be a maximal subforest of the union of the S_u, where the union ranges over all upper neighbours u of t. We recursively build a maximal subset U of $V(T)$ such that if $a, b \in U$, then Q_a and Q_b are vertex-disjoint. In this construction, we first add the nodes of T with smaller distance from the root. This ensures by the "Moreover"-part of Remark 6.6 that U contains infinitely many nodes of each ray of T.

Let S' be the union of those Q_t with $t \in U$. We obtain S by extending S' to a spanning tree of G, and rooting it at some $v \in V(S)$ arbitrarily. By the Star-Comb-Lemma [14, Section 8], each spanning tree of G contains for each topological end ω a ray belonging to ω.

Thus it remains to show that S does not contain two disjoint rays R_1 and R_2 that both belong to the same topological end ω of G. Suppose there are such R_1, R_2 and ω. Let $t_1t_2\dots$ be the ray of T in which ω lives. Let n

be so large that both R_1 and R_2 meet P_{t_n}. Then for each $m \geq n$, the set S_{t_m} contains a path joining R_1 and R_2. Thus the set $Q_{t_{m-1}}$ contains such a path. Since $Q_{t_{m-1}} \subseteq S$ for infinitely many m, the tree S contains a cycle, which is the desired contradiction. $\square$

Remark 6.14. In the above proof of Corollary 2.6 in the application of the Star-Comb-Lemma we used the property that topological ends are not dominated by vertices.

However, with a little bit more care, one can show more generally that if a graph has a tree-decomposition such that the ends of the decomposition tree define precisely a set Ψ of vertex-ends of the graph, then this graph has a spanning tree that is end-faithful for that set Ψ.

To see that we show that if a graph G has such a tree-decomposition $(T, P_t | t \in V(T))$, then it has a connected subgraph G' with the same vertex set such that all vertex-ends in the set Ψ are vertex-ends of G' that are topological.

It is fairly easy to see that we may assume that the tree-decomposition $(T, P_t | t \in V(T))$ has the following additional properties.

1. By Q_t we denote the union of the part P_t with all parts P_u, where u is above t in the decomposition tree, without the part P_s; here s is the downward-neighbour of t and t is not the root. Then the graph Q_t is connected.

2. Every part P_t contains a finite connected set C_t such that $C_t \subseteq P_t \setminus P_s$ and every vertex of the separator $V(P_t) \cap V(P_s)$ is in the neighbourhood of C_t; here s and t are as in the first property.

3. Let $s \leq t \leq u$ such that $st, tu \in E(T)$. Then the separator $V(P_t) \cap V(P_u)$ is not a subset of $V(P_s) \cap V(P_t)$.

For a node t different from the root, let K_t be the union of P_t with all sets C_u, where u is an upward-neighbour of t, without the part P_s, where s is the down-ward neighbour of t. Since the set Q_t is connected and no separator $V(P_t) \cap V(P_u)$ is a subset of P_s by the third property, the graph K_t must be connected. If t is the root, we define K_t the same but without removing a part P_s; this graph is connected as the graph G is connected.

We define G' to be the union of the connected subgraphs K_t. This graph is connected. It is straightforward to check that every vertex-end of Ψ is a topological end of G'. So we can apply Corollary 2.6 to the graph G' to deduce that the graph G has a spanning tree that is end-faithful for the set Ψ.

7 Concluding Remarks

We have shown that any graph has a tree-decomposition of finite adhesion that distinguishes its topological ends. It is natural to ask whether such a statement is true if we replace 'topological ends' by some other classes of vertex-ends.

Let us be more precise. A class $\mathcal{C}$ of vertex-ends is *tree-distinguishable* if every graph has a tree-decomposition of finite adhesion that distinguishes any two vertex-ends that are in $\mathcal{C}$. We would like to know which natural classes $\mathcal{C}$ of vertex-ends are tree-distinguishable?

As demonstrated in Example 3.1 the class of all vertex-ends is not tree-distinguishable. A class that has received a lot of attention in the literature, see for example [25], is the class of 'k-thin' vertex-ends; here, given a natural number k, a vertex-end ω is *k-thin* if the number k_1 of vertices dominating ω and the cardinality k_2 of any family of vertex-disjoint rays belonging to ω sum up to at most k, that is $k_1 + k_2 \leq k$. For example, thin ends are k-thin for every sufficiently large value of k; indeed vertex-ends dominated by infinitely many vertices have infinitely many vertex-disjoint rays belonging to that end, see [15]. Although the class of thin vertex-ends is not tree-distinguishable by Example 3.3, the class of k-thin vertex-ends is; this is well-known and also follows from Theorem 7.1 below.

For general graphs the class of k-thin vertex-ends and topological ends are not subsets of one another. Is there a natural tree-distinguishable class that contains both of them?

Yes, there is such a class and the proof that it is tree-distinguishable is an easy application of our main theorem, as follows. Given a natural number k, a vertex-end is *k-dominated* if it is dominated by at most k vertices. For example, the 0-dominated vertex-ends are the topological ends. Clearly every k-thin vertex-end is k-dominated. The following extension of Theorem 1 implies that the class of k-dominated vertex-ends is tree-distinguishable for any fixed k.

Theorem 7.1. *For any fixed natural number k, every graph has a tree-decomposition $(T, \mathcal{V})$ of finite adhesion such that the ends of T define precisely the k-dominated vertex-ends of G.*

Proof. We shall prove Theorem 7.1 by induction on k. All trees in this proof are rooted; and we denote their root by r. Along that induction we shall prove the following property: if a vertex-end lives in a part P_t of $(T, \mathcal{V})$, then it is dominated by $k+1$ vertices contained in the separator $V(P_t) \cap V(P_s)$, where s is the neighbour of t in the tree T that is nearer to the root. (In particular, no vertex-end lives in the root part P_r, which may be assumed to consist of finitely many vertices).

We remark that in our proof of Theorem 1 we could construct the tree-decomposition $(T, \mathcal{V})$ such that if a vertex-end lives in a part P_t, then it is dominated by at least one vertex contained in the separator $V(P_t) \cap V(P_s)$ (with s as above). Indeed, we just have to use the variant of Lemma 6.1, where we replace 'each topological end' by 'each vertex-end not dominated by a vertex of W'. This variant is deduced with the same proof from Lemma 6.7 except that we replace 'topological end' by 'vertex-end not dominated by a vertex of the set W'. During this proof we assume that Theorem 1 includes this additional statement.

So the base case $k = 0$ follows from Theorem 1.

Now assume that we already have a suitable tree-decomposition $(T, \mathcal{V}) = (T, P_t | t \in V(T))$ for k. We take each torso of a part of that tree-decomposition and delete the separator $V(P_t) \cap V(P_s)$. Call the resulting graph H_t. Now we apply Theorem 1 to the graph H_t. Call this tree-decomposition $(T[t], \mathcal{V}[t])$. We obtain a tree-decomposition of the part P_t by adding the separator $V(P_t) \cap V(P_s)$ to all parts of the tree-decomposition $(T[t], \mathcal{V}[t])$. Call that tree-decomposition $(T'[t], \mathcal{V}'[t])$.

We obtain a suitable tree-decomposition from $(T, \mathcal{V})$ by replacing each part P_t by the tree-decomposition $(T'[t], \mathcal{V}'[t])$. This is well-defined as each separator of $(T, \mathcal{V})$ is a complete subgraph of H_t. Indeed, then we can attach the parts of $(T, \mathcal{V})$ above P_t at the unique part of $(T'[t], \mathcal{V}'[t])$ nearest to the root that includes the corresponding separator. It is straightforward to check that this tree-decomposition has the desired property. $\square$

In this paper we considered various classes of vertex-ends. In Figure 12, we depict the inclusion-relations that hold between the classes of vertex-ends considered in this paper. In Figure 13 we summarise which classes of vertex-ends are tree-distinguishable and which classes have end-faithful spanning trees. We recall that countable graphs have normal spanning trees and hence end-faithful spanning trees for vertex-ends. Hence the questions of this paper are of particular interest for uncountable graphs.

All positive results of Figure 13 are proved in this paper. The counterexamples corresponding to the cross in the bottom right corner were constructed by Seymour and Thomas and Thomassen as mentioned in the introduction. The other two crosses are derived from Example 3.1 and Example 3.3. The question mark in Figure 13 corresponds to the following question.

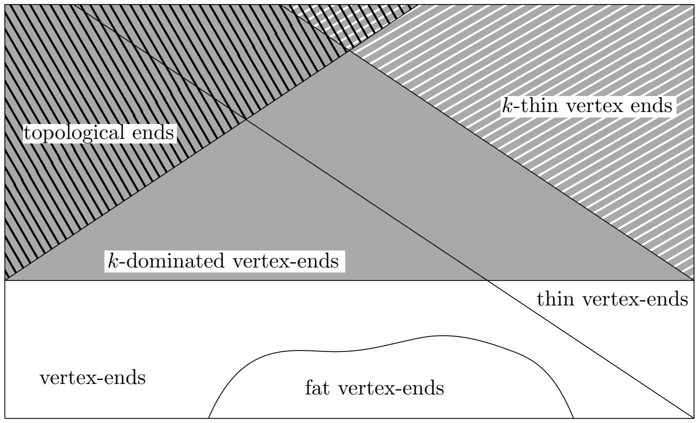

Figure 12: The classes of vertex-ends considered in this paper.

	tree-distinguishable	end-faithful spanning tree
k-thin	✓	✓
topological	✓	✓
k-dominated	✓	✓
thin	×	?
all vertex-ends	×	×

Figure 13: We put a tick in an entry of this table if the corresponding class of vertex-ends is tree-distinguishable or there is an end-faithful spanning tree for that class, respectively. If it is false we put a cross. In the one case where it is open, we put a question mark.

Question 7.2. *Does every graph have an end-faithful spanning tree for the thin vertex-ends?*

The strengthening of Question 7.2 with ‘thin’ replaced by the class of vertex-ends that are dominated by finitely (or more generally: countably) many vertices is also open. Since this class contains the topological ends, this possible strengthening implies Corollary 2.6.

However, by Example 3.3 the class of thin vertex-ends is not tree-distinguishable. Thus, if the answer to Question 7.2 was ‘yes’, the obvious strategy suggested by Remark 6.14 directly via tree-decompositions defining precisely the thin ends cannot succeed.

In the introduction of this paper we claimed that we repaired Halin’s Conjecture. In the presence of Question 7.2 this might deserve some further justification. Given the counterexamples against Halin’s original conjecture, we explain the subtle difference between the following two questions.

1. How can Halin’s Conjecture be repaired?

2. What is the largest possible natural subclass of the vertex-ends for which the weakening of Halin’s Conjecture is true?

The second question is still open. Candidates for that subclass are the k-dominated vertex-ends, or more generally the finitely or countably dominated ones. These subclasses are vertex-ends with some finiteness or countability assumption. Unlike for the topological ends, it would not have been natural to ask the weakening of Halin’s Conjecture for those subclasses in 1964. Hence these subclasses can hardly give answers to question one. The thin vertex-ends are also vertex-ends together with some finiteness assumption and might have served as a solution to question one if Diestel’s original problem was true. Since the original problem is not true (and in fact can be repaired for the topological ends), the author is convinced that Corollary 2.6 is the most natural way to repair Halin’s Conjecture.

Acknowledgement

I thank two anonymous referees whose valuable comments improved this paper. Furthermore I thank Nathan Bowler for pointing out an error in an earlier version.

References

[1] Eli Berger and Henning Bruhn. Eulerian edge sets in locally finite graphs. *Combinatorica*, 31(1):21–38, 2011.

[2] N. Bowler and J. Carmesin. Infinite matroids and determinacy of games. Preprint 2013, current version available at http://arxiv.org/abs/1301.5980.

[3] H. Bruhn and M. Stein. On end degrees and infinite circuits in locally finite graphs. *Combinatorica*, 27:269–291, 2007.

[4] C. Carathéodory. Über die Begrenzung einfach zusammenhängender Gebiete. *Math. Ann.*, 73(3):323–370, 1913.

[5] J. Carmesin. All graphs have tree-decompositions displaying their topological ends. Preprint 2015, available at http://arxiv.org/pdf/1409.6640v4.

[6] J. Carmesin. On the end structure of infinite graphs. Preprint 2014, available at http://arxiv.org/pdf/1409.6640v1.

[7] J. Carmesin. A short proof that every finite graph has a tree-decomposition displaying its tangles. *European J. Combin.*, 58:61–65, 2016.

[8] J. Carmesin. Topological cycle matroids of infinite graphs. *European J. Combin*, 60:135–150, 2017.

[9] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. k-blocks: a connectivity invariant for graphs. *SIAM J. Discrete Math.*, 28(4):1876–1891, 2014.

[10] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree decompositions of finite graphs I: Existence and algorithms. *J. Combin. Theory (Series B)*, pages 1–24, 2016.

[11] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. Connectivity and tree-structure in finite graphs. *Combinatorica*, 34:11–46, 2014.

[12] R. Diestel. Locally finite graphs with ends: a topological approach. http://arxiv.org/abs/0912.4213.

[13] R. Diestel. The end structure of a graph: Recent results and open problems. *Disc. Math.*, 100:313–327, 1992.

[14] R. Diestel. *Graph Theory* (4th edition). Springer-Verlag, 2010. Electronic edition available at:
`http://diestel-graph-theory.com/index.html`.

[15] R. Diestel. *Graph Theory* (5th edition). Springer-Verlag, 2016. Electronic edition available at:
`http://diestel-graph-theory.com/index.html`.

[16] R. Diestel. Ends and tangles. *Abhandlungen Math. Sem. Univ. Hamburg*, 2017.

[17] R. Diestel and D. Kühn. Graph-theoretical versus topological ends of graphs. *J. Combin. Theory (Series B)*, 87:197–206, 2003.

[18] R. Diestel and D. Kühn. On infinite cycles I. *Combinatorica*, 24:68–89, 2004.

[19] R. Diestel and D. Kühn. On infinite cycles II. *Combinatorica*, 24:91–116, 2004.

[20] H. Freudenthal. Über die Enden topologischer Räume und Gruppen. *Math. Zeitschr.*, 33:692–713, 1931.

[21] Hans Freudenthal. Über die Enden diskreter Räume und Gruppen. *Comment. Math. Helv.*, 17:1–38, 1945.

[22] A. Georgakopoulos. Infinite Hamilton cycles in squares of locally finite graphs. *Advances in Mathematics*, 220:670–705, 2009.

[23] R. Halin. Über unendliche Wege in Graphen. *Math. Annalen*, 157:125–137, 1964.

[24] R. Halin. Lattices of cuts in graphs. *Abh. Math. Sem. Univ. Hamburg*, 61:217–230, 1991.

[25] M. Hamann and J. Pott. Transitivity conditions in infinite graphs. *Combinatorica*, 32:649–688, 2012.

[26] B. Hughes and A. Ranicki. *Ends of complexes.* Cambridge Tracts in Mathematics 123, Cambridge Univ. Press, 1996.

[27] P. Seymour and R. Thomas. An end-faithful spanning tree counterexample. *Proc. Amer. Math. Soc.*, 113:1163–1171, 1991, no. 4.

[28] M. Stein. Arboriticity and tree-packing in locally finite graphs. *J. Combin. Theory (Series B)*, 96:302–312, 2006.

[29] R. Thomas. Well-quasi-ordering infinite graphs with forbidden finite planar minor. *Trans. Amer. Math. Soc.*, 312:279–313, 1989, no. 1.

[30] C. Thomassen. Infinite connected graphs with no end-preserving spanning trees. *J. Combin. Theory (Series B)*, 54:322–324, 1992.