

Refining a Tree-Decomposition which Distinguishes Tangles

Joshua Erde

March 6, 2017

Abstract

Robertson and Seymour introduced tangles of order k as objects representing highly connected parts of a graph and showed that every graph admits a tree-decomposition of adhesion $< k$ in which each tangle of order k is contained in a different part. Recently, Carmesin, Diestel, Hamann and Hundertmark showed that such a tree-decomposition can be constructed in a canonical way, which makes it invariant under automorphisms of the graph. These canonical tree-decompositions necessarily have parts which contain no tangle of order k , which we call inessential. Diestel asked what could be said about the structure of the inessential parts. In this paper we show that the torsos of the inessential parts in these tree-decompositions have branch-width $< k$, allowing us to further refine the canonical tree-decompositions, and also show that a similar result holds for k -blocks. We also use our methods to further refine the essential parts in such a tree-decomposition in a similar fashion.

1 Introduction

A classical notion in graph theory is that of the block-cut vertex tree of a graph. It tells us that if we consider the maximal 2-connected components of a connected graph G then they are arranged in a ‘tree-like’ manner, separated by the cut vertices of G . A result of Tutte’s [11] says that we can decompose any 2-connected graph in a similar way. Broadly, it says that every 2-connected graph can be decomposed in a ‘tree-like’ manner, so that the parts are separated by vertex sets of size at most 2, and every part, together with the edges in the separators adjacent to it, is either 3-connected or a cycle. We call the union of a part and the edges in the separators adjacent to it the *torso* of the part. In contrast to the first example not every part, or even torso, of this decomposition is 3-connected, and indeed it is easy to show that not every 2-connected graph can be decomposed in this way such that every torso is 3-connected.

It has long been an open problem how best to extend these results for general k , the aim being to decompose a $(k - 1)$ -connected graph into its ‘ k -connected components’, where the precise meaning of what these ‘ k -connected components’ should be considered to be has varied. Tutte’s example shows us that there may be parts of this decomposition which are not highly connected, but rather play a structural role in the graph of linking the highly connected parts together, and further that the highly connected parts of the decomposition may not correspond exactly to k -connected subgraphs.

Whereas initially these ‘ k -connected components’ were considered as concrete structures in the graph itself, Robertson and Seymour [9] radically re-interpreted them as *tangles of order k* , which for brevity we will refer to as *k -tangles*¹. Instead of being defined in terms of the edges and vertices of a graph, these objects were defined in terms of structures on the set of low-order separations of a graph.

Robertson and Seymour showed that, given any set of distinct k -tangles T_1, T_2, \dots, T_n in a graph G , there is a tree-decomposition (T, \mathcal{V}) of G with precisely n parts in which the orientations induced by the tangles T_i on $E(T)$ each have distinct sinks, where we say the tangle is *contained* in this sink. We say that such a tree-decomposition *distinguishes* the tangles T_1, T_2, \dots, T_n . They showed further that these tree-decompositions can be chosen so that the separators between the parts are in some way minimal with respect to the tangles considered. We say that such a tree-decomposition distinguishes the k -tangles *efficiently*. If we call the largest size of a separator in a tree-decomposition the *adhesion* of the tree-decomposition, then in particular their result implies the following:

¹Precise definitions of many of the terms in the introduction will be postponed until Section 2, where all the necessary background material will be introduced.

Theorem 1 (Robertson and Seymour [9]). *For every graph G and $k \geq 2$ there exists a tree-decomposition (T, \mathcal{V}) of G of adhesion $< k$ which distinguishes the set of k -tangles in G efficiently.*

More recently Carmesin, Diestel, Hamann and Hundertmark [1] described a family of algorithms that can be used to build tree-decompositions which distinguish the set of k -tangles in a graph and are *canonical*, that is, they are invariant under every automorphism of the graph.

Just as in Tutte's theorem, where there were parts of the tree-decomposition whose torsos were not 3-connected, it is easy to show that the tree-decompositions formed in [1] must contain parts which do not contain any k -tangle. Since the general motivation for these tree-decompositions is to decompose the graph into its ' k -connected components' in a way that displays the global structure of the graph, it is natural to ask further questions about the structure of these tree-decompositions. In [2] Carmesin et al. analysed the structure of the trees that the various algorithms given in [1] produced. One particular question that was asked is what can be said about the structure of the parts which do not contain a k -tangle. We will call the parts of a tree-decomposition that contain a k -tangle *essential*, and those that do not *inessential*.

For example, if the whole graph contains no k -tangle, then these canonical tree-decompositions tell us nothing about the graph, as they consist of just one inessential part. However there are theorems which describe the structure of a graph which contains no k -tangle. In the same paper where they introduced the concept of tangles, Robertson and Seymour [9] showed that a graph which contains no k -tangle has branch-width $< k$, and in fact that the converse is also true, a graph with branch-width $\geq k$ contains a k -tangle. Having branch-width $< k$ can be rephrased in terms of the existence of a certain type of tree-decomposition (See e.g. [8]). A nice property of these tree-decompositions is that each of the parts is in some sense 'too small' to contain a k -tangle. In this way these tree-decompositions witness that a graph has no k -tangle by splitting the graph into a number of parts, each of which cannot contain a k -tangle and similarly a k -tangle witnesses that a graph does not have such a tree-decomposition.

A natural question to then ask is, do the inessential parts in the tree-decompositions from [1] admit tree-decompositions of the same form, into parts which are too small to contain a k -tangle? If so we might hope to refine these canonical tree-decompositions by decomposing further the inessential parts. By combining these decompositions we would get an overall tree-decomposition of G consisting of some essential parts, each containing a k -tangle in G , and some inessential parts, each of which is 'small' enough to witness the fact that no k -tangle is contained in that part.

We first note that we cannot hope for these refinements to also be canonical. For example consider a graph formed by taking a large cycle C and adjoining to each edge a large complete graph K_n . Then a canonical tree-decomposition which distinguishes the 3-tangles in this graph will contain the cycle C as an inessential part. However there is no canonical tree-decomposition of C with branch-width < 3 . Indeed, if such a tree-decomposition contained any of the 2-separations of C as an adhesion set then, since all the rotations of C lie in the automorphism group of G , every rotation of this separation must appear as an adhesion set. However these separations cannot all appear as the adhesion sets in any tree-decomposition, as every pair of vertices in a 2-separation of C are themselves separated by some rotation of that separation.

If we drop the restriction that the refinement be canonical then, at first glance, it might seem like there should clearly be such a refinement. If there is no k -tangle contained in a part V_t in a tree-decomposition, (T, \mathcal{V}) , then by the theorem of Robertson and Seymour there should be a tree-decomposition of that part with branch-width $< k$. However there is a problem with this naive approach, in that we have no guarantee that we can insert the tree-decomposition of this part into the existing tree-decomposition. In particular it could be the case that this tree-decomposition splits up the separators of the part V_t in (T, \mathcal{V}) . One way to avoid this problem is to instead consider the torso of the part V_t . If we have a tree-decomposition of the torso we can insert it into the original tree-decomposition, but it is not clear that adding these extra edges can not increase the branch-width of the part. In fact it is easy to find examples where choosing a bad canonical tree-decomposition to distinguish the set of k -tangles in a graph results in inessential parts whose torsos have branch-width $\geq k$.

For example consider the following graph: We start with the union of three large complete graphs, K_{N_1} , K_{N_2} and K_{N_3} , for $N_1, N_2, N_3 \gg k$. We pick a set of $(k-1)/2$ vertices from each graph, which we denote by X_1 , X_2 and X_3 respectively, and join each of these sets completely to a new vertex x . It is a simple check that there are three k -tangles in this graph, corresponding to the three large complete subgraphs. However, consider the following tree-decomposition of the graph into four parts $K_{N_1} \cup X_2$,

$K_{N_2} \cup X_3$, $K_{N_3} \cup X_1$ and $X_1 \cup X_2 \cup X_3 \cup \{x\}$. This is a tree-decomposition which distinguishes the k -tangles in the graph, and the part $X_1 \cup X_2 \cup X_3 \cup \{x\}$ is inessential. However the torso of this middle part is a complete graph of order $3(k-1)/2 + 1$, which can be seen to have branch-width $\geq k$.




Figure 1: A graph with a bad tangle-distinguishing tree-decomposition.

We will show that, for the canonical tree-decompositions of Carmesin et al, the torsos of the inessential parts all have branch-width $< k$ and so it is possible to decompose the torsos of the inessential parts in this way.

Theorem 2. *For every graph G and $k \geq 3$ there exists a canonical tree-decomposition (T, \mathcal{V}) of G of adhesion $< k$ such that*

- (T, \mathcal{V}) distinguishes the set of k -tangles in G efficiently;
- The torso of every inessential part has branch-width $< k$.

More recently another potential candidate for these ‘ k -connected components’ has been considered in the literature, called k -blocks. We say that a set of at least k vertices in a graph is $(< k)$ -inseparable if no set of $< k$ vertices can separate any two of the vertices. A k -block is a maximal $(< k)$ -inseparable set of vertices. These k -blocks differ from subgraphs which are k -connected in the classical sense in that their connectivity is measured in the ambient graph rather than the subgraph itself. For example if we take a large independent set, I , and join each pair of vertices in I by k vertex disjoint paths, then I is a k -block, even though as a subgraph it is independent. Carmesin, Diestel, Hundertmark and Stein [3] showed that, for any graph G , there is a canonical tree-decomposition which distinguishes the set of k -blocks. The work of Carmesin et al [1] extended the results of [3] to more general types of highly connected substructures in graphs, and these results have been extended further by Diestel, Hundertmark and Lemanczyk [5] to more general combinatorial structures, such as matroids.

As before, these tree-decompositions will have some parts which are *essential*, that is they contain a k -block, and some parts which are *inessential*, and it is natural to ask about the structure of these parts. Recently, Diestel, Eberenz and Erde [4] proved a duality theorem for k -blocks, analogous to the tangle/branch-width duality of Robertson and Seymour. The result implies that a graph contains a k -block if and only if it does not admit a tree-decomposition of *block-width* $< k$, where as before, every part in a tree-decomposition of block-width $< k$ is in some sense ‘too small’ to contain a k -block. We also show a corresponding result for blocks.

Theorem 3. *For every graph G and $k \geq 3$ there exists a canonical tree-decomposition (T, \mathcal{V}) of G of adhesion $< k$ such that*

- (T, \mathcal{V}) distinguishes the set of k -blocks in G efficiently;
- The torso of every inessential part has block-width $< k$.

The main result in this paper, of which Theorems 2 and 3 are corollaries, is a lemma that gives sufficient conditions on the separators of an inessential part in a distinguishing tree-decomposition for the torso to have small width. These conditions seem quite natural and reasonable, in particular they are satisfied by every part of the canonical tangle/block-distinguishing tree-decompositions constructed by Carmesin et al.

In some sense the canonical tangle-distinguishing tree-decompositions tell us most about the structure of the graph when the essential parts correspond closely to the tangles inside them. For example consider the following two graphs, firstly two K_N s overlapping in $k-1$ vertices and secondly two $K_{3k/2}$ s each with a long path attached, of length $N' = N - 3k/2$, overlapping in a similar way, see Figure 2.




Figure 2: Two graphs with the same canonical k -tangle-distinguishing tree-decomposition.

Since the tangle-distinguishing tree-decompositions of Carmesin et al. only use essential separations, that is separations which distinguish some pair of k -tangles, they will construct the same tree-decomposition for both of these graphs, with just two parts of size N . However in the second example a more sensible tree-decomposition would further split up the long paths. This could be done in a way to maintain the property that the inessential parts have small branch-width, and by separating these inessential parts from the essential part we have more precisely exhibited the structure of the graph. We will also apply our methods to the problem of further refining the essential parts of these tree-decompositions.

In Section 2 we introduce the background material necessary for our proof and in Section 3 we prove our central lemma and deduce the main results in the paper. In Section 4 we discuss how our methods can also be used to further refine the essential parts of a tree-decomposition.

2 Background material

2.1 Separation systems and tree-decompositions

A *separation* of a graph G is a set $\{A, B\}$ of subsets of $V(G)$ such that $A \cup B = V$ and there is no edge of G between $A \setminus B$ and $B \setminus A$. There are two *oriented separations* associated with a separation, (A, B) and (B, A) . Informally we think of (A, B) as *pointing towards* B and *away from* A . We can define a partial ordering on the set of oriented separations of G by

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

The *inverse* of an oriented separation (A, B) is the separation (B, A) , and we note that mapping every oriented separation to its inverse is an involution which reverses the partial ordering.

In [7] Diestel and Oum generalised these properties of separations of graphs and worked in a more abstract setting. They defined a *separation system* $(\vec{S}, \leq, *)$ to be a partially ordered set \vec{S} with an order reversing involution, $*$. The elements of \vec{S} are called *oriented separations*. Often a given element of \vec{S} is denoted by \vec{s} , in which case its inverse \vec{s}^* will be denoted by \overleftarrow{s} , and vice versa. Since $*$ is ordering reversing we have that, for all $\vec{r}, \vec{s} \in \vec{S}$,

$$\vec{r} \leq \vec{s} \text{ if and only if } \overleftarrow{r} \geq \overleftarrow{s}.$$

A *separation* is a set of the form $\{\vec{s}, \overleftarrow{s}\}$, and will be denoted by simply s . The two elements \vec{s} and \overleftarrow{s} are the *orientations* of s . The set of all such pairs $\{\vec{s}, \overleftarrow{s}\} \subseteq \vec{S}$ will be denoted by S . If $\vec{s} = \overleftarrow{s}$

we say s is *degenerate*. Conversely, given a set $S' \subseteq S$ of separations we write $\vec{S}' := \bigcup S'$ for the set of all orientations of its elements. With the ordering and involution induced from \vec{S} , this will form a separation system. When we refer to an oriented separation in a context where the notation explicitly indicates orientation, such as \vec{s} or (A, B) , we will usually suppress the prefix “oriented” to improve the flow of the paper.

Given a separation of a graph $\{A, B\}$ we can identify it with the pair $\{(A, B), (B, A)\}$ and in this way any set of separations in a graph which is closed under taking inverses forms a separation system. We will work within the framework developed in [7] since we will need to use directly some results proved in this abstract setting, but also because our results are most easily expressible in this framework. An effort has been made to state the results in the widest generality, so as to be applicable in the broadest sense, however we will always have in mind the motivating example of separation systems which arise as sets of separations in a graph, and so a reader will not lose too much by thinking about these separation systems solely in those terms.

The *separator* of a separation $\vec{s} = (A, B)$ in a graph is the intersection $A \cap B$ and the *order* of a separation, $|\vec{s}| = \text{ord}(A, B)$, is the cardinality of the separator $|A \cap B|$. Note that if $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$ are separations then so are the *corner separations* $\vec{r} \vee \vec{s} := (A \cup C, B \cap D)$ and $\vec{r} \wedge \vec{s} := (A \cap C, B \cup D)$ and the orders of these separations satisfy the equality

$$|\vec{r} \vee \vec{s}| + |\vec{r} \wedge \vec{s}| = |\vec{r}| + |\vec{s}|.$$

Hence the order function is a submodular function on the set of separations of a graph, and we note also that it is clearly symmetric.




Figure 3: Two separations (A, B) and (C, D) with the corner separation $(A \cup C, B \cap D)$ marked.

For abstract separations systems, if there exists binary operations \vee and \wedge on \vec{S} such that $\vec{r} \vee \vec{s}$ is the supremum and $\vec{r} \wedge \vec{s}$ is the infimum of \vec{r} and \vec{s} then we call $(\vec{S}, \leq, *, \vee, \wedge)$ a *universe* of (oriented) separations, and we call any real, non-negative, symmetric and submodular function on a universe an *order function*.

Two separations r and s are *nested* if they have \leq -comparable orientations. Two oriented separations \vec{r} and \vec{s} are *nested* if r and s are nested². If \vec{r} and \vec{s} are not nested we say that the two separations *cross*. A set of separations S is *nested* if every pair of separations in S is nested, and a separation s is *nested* with a set of separations S if $S \cup \{s\}$ is nested.

A separation $\vec{r} \in \vec{S}$ is *trivial* in \vec{S} , and \overleftarrow{r} is *co-trivial*, if there exist an $s \in S$ such that $\vec{r} < \vec{s}$ and $\vec{r} < \overleftarrow{s}$. Note that if \vec{r} is trivial, witnessed by some s , then, since the involution is order reversing, we have that $\vec{r} < \vec{s} < \overleftarrow{r}$. So, in particular, \overleftarrow{r} cannot also be trivial. Separations \vec{s} such that $\vec{s} \leq \overleftarrow{s}$, trivial or not, will be called *small* and their inverses *co-small*.

In the case of separations of a graph, it is a simple check that the small separations are precisely those of the form (A, V) . Furthermore the trivial separations can be characterised as those of the form (A, V)

²In general we will use terms defined for separations informally for oriented separations when the meaning is clear, and vice versa

such that $A \subseteq C \cap D$ for some separation (C, D) such that $\{C, D\} \neq \{A, B\}$. Finally we note that there is only one degenerate separation in a graph, (V, V) .

A *tree-decomposition* of a graph G is a pair (T, \mathcal{V}) consisting of a tree T and family $\mathcal{V} = (V_t)_{t \in T}$ of vertex sets $V_t \subseteq V(G)$, one for each vertex $t \in T$ such that:

- $V(G) = \bigcup_{t \in T} V_t$;
- for every edge $e \in G$ there exists some $t \in T$ such that $e \in G[V_t]$;
- $V_{t_1} \cap V_{t_2} \subseteq V_{t_3}$ whenever t_3 lies on the $t_1 - t_2$ path in T .

The sets V_t in a tree-decomposition are its *parts* and the sets $V_t \cap V_{t'}$ such that (t, t') is an edge of T are the *adhesion sets*. The *torso* of a part \overline{V}_t is the union of that part together with the completion of the adhesion sets adjacent to that part, that is

$$\overline{V}_t = G|_{V_t} \cup \bigcup_{(t, t') \in T} K_{V_t \cap V_{t'}}.$$

The *width* of a tree-decomposition is $\max\{|V_t| - 1 : \text{such that } t \in T\}$, and the *adhesion* is the size of the largest adhesion set. Deleting an oriented edge $e = (t_1, t_2) \in \vec{E}(T)$ divides $T - e$ into two components $T_1 \ni t_1$ and $T_2 \ni t_2$. Then $(\bigcup_{t \in T_1} V_t, \bigcup_{t \in T_2} V_t)$ can be seen to be a separation of G with separator $V_{t_1} \cap V_{t_2}$. We say that the edge e *induces* this separation. Given a tree-decomposition (T, \mathcal{V}) it is easy to check that the set of separations induced by the edges of T form a nested separation system. Conversely it was shown in [3] that every nested separation system is induced by some tree-decomposition, and so in a sense these two concepts can be thought of as equivalent.

We say that a nested set of separations \mathcal{N}' *refines* a nested set of separations \mathcal{N} if $\mathcal{N}' \supseteq \mathcal{N}$, and similarly a tree-decomposition (T', \mathcal{V}') *refines* a tree-decomposition (T, \mathcal{V}) if the set of separations induced by the edges of T' refines the corresponding set of separations for T .

2.2 Duality of tree-decompositions

There are a number of theorems that assert a duality between certain structurally ‘large’ objects in a graph and an overall tree structure. For example a graph has small tree-width if and only if it contains no large order bramble [10]. In [7] a general theory of duality, in terms of separation systems, was developed which implied many of the existing theorems. Following on from the notion of tangles in graph minor theory [9] these large objects were described as orientations of separations systems avoiding certain forbidden subsets.

An *orientation* of a set of separations S is a subset $O \subseteq \vec{S}$ which for each $s \in S$ contains exactly one of its orientations \vec{s} or \overleftarrow{s} . A *partial orientation* of S is an orientation of some subset of S , and we say that an orientation O *extends* a partial orientation P if $P \subseteq O$.

In our context we will think of an orientation O on some set of graph separations as choosing a side of each separation $s = \{A, B\}$ to designate as large. For example given a graph G and the set S of all separations of the graph G , we denote by

$$\vec{S}_k = \{\vec{s} \in \vec{S} : |\vec{s}| < k\},$$

the set of all orientations of order less than k . If there is a large clique (of size $\geq k$) in G then for every $s = \{A, B\} \in S_k$ we have that the clique is contained entirely in A or B . So this clique defines an orientation of S_k by picking, for each $\{A, B\} \in S_k$ the orientated separation such that the clique is contained in second set in the pair.

We call an orientation O of a set of separations S *consistent* if whenever we have distinct r and s such that $\vec{r} < \vec{s}$, O does not contain both \overleftarrow{r} and \vec{s} . Note that a consistent orientation must contain all trivial separations \vec{r} , since if $\vec{r} < \vec{s}$ and $\vec{r} < \overleftarrow{s}$ then, whichever orientation of s is contained in O would be inconsistent with \overleftarrow{r} .

Given a set of subsets $\mathcal{F} \subseteq 2^{\vec{S}}$ we say that an orientation O is \mathcal{F} -*avoiding* if there is no $F \in \mathcal{F}$ such that $F \subseteq O$. So for example an orientation is consistent if it avoids $\mathcal{F} = \{\{\overleftarrow{r}, \vec{s}\} : r \neq s, \vec{r} < \vec{s}\}$. In general we will define the ‘large’ objects we consider by the collection \mathcal{F} of subsets they avoid. For

example a k -tangle in a graph G can easily be seen to be equivalent to an orientation of S_k which avoids the set of triples

$$\mathcal{T}_k = \{(A_1, B_1), (A_2, B_2), (A_3, B_3)\} \subseteq \vec{S}_k : \bigcup_{i=1}^3 G[A_i] = G\}.$$

(Where the three separations need not be distinct). That is, a tangle is an orientation such that no three small sides cover the entire graph, it is a simple check that any such orientation must in fact also be consistent. We say that a consistent orientation which avoids a set \mathcal{F} is an \mathcal{F} -tangle.

Given a set $\mathcal{F} \subseteq 2^{\vec{S}}$, an S -tree over \mathcal{F} is a pair (T, α) , of a tree T with at least one edge and a function $\alpha : \vec{E}(T) \rightarrow \vec{S}$ from the set

$$\vec{E}(T) := \{(x, y) : \{x, y\} \in E(T)\}$$

of orientations of it's edges to \vec{S} such that:

- For each edge $(t_1, t_2) \in \vec{E}(T)$, if $\alpha(t_1, t_2) = \vec{s}$ then $\alpha(t_2, t_1) = \overleftarrow{s}$;
- For each vertex $t \in T$, the set $\{\alpha(t', t) : (t', t) \in \vec{E}(T)\}$ is in \mathcal{F} ;

For any leaf vertex $w \in T$ which is adjacent to some vertex $u \in T$ we call the separation $\vec{s} = \alpha(w, u)$ a *leaf separation* of (T, α) . A particularly interesting class of such trees is when the set \mathcal{F} is chosen to consist of stars. A set of non-degenerate oriented separations σ is called a *star* if $\vec{r} \leq \overleftarrow{s}$ for all distinct $\vec{r}, \overleftarrow{s} \in \sigma$. In what follows, if we refer to an S -tree without reference to a specific family \mathcal{F} of stars, it can be assumed to be over the set of *all* stars in $2^{\vec{S}}$. We say that an S -tree over \mathcal{F} is *irredundant* if there is no $t \in T$ with two neighbours, t' and t'' such that $\alpha(t, t') = \alpha(t, t'')$. If (T, α) is an irredundant S -tree over a set of stars \mathcal{F} , then it is easy to verify that the map α preserves the natural ordering on $\vec{E}(T)$, defined by letting $(s, t) \leq (u, v)$ if the unique path in T between those edges starts at s and ends at v (see [[7], Lemma 2.2]).

Given an irredundant S -tree (T, α) over a set of stars and an orientation O of S , O induces an orientation of the edges of T , which will necessarily contain a sink vertex. If the orientation O is consistent then this sink vertex, which we will denote by t , will be unique. We say that O is *contained in* t . If $S = S_k$ for some graph G , we have that (T, α) defines some tree-decomposition (T, \mathcal{V}) of G , and we say that O is *contained in the part* V_t . So, each \mathcal{F} -tangle of S must live in some vertex of every such S -tree, and by definition this vertex give rise to a star of separations in \mathcal{F} . In this way, each of the vertices in an S -tree over \mathcal{F} (and each of the parts in the corresponding tree-decomposition when one exists) is ‘too small’ to contain an \mathcal{F} -tangle.

Suppose we have a separation \vec{r} which is neither trivial nor degenerate. In applications \vec{r} will be a leaf separation in some irredundant S -tree over a set \mathcal{F} of stars. Given some $\vec{s} \geq \vec{r}$, it will be useful to have a procedure to ‘shift’ the S -tree (T, α) in which \vec{r} is a leaf separation to a new S -tree (T, α') such that \vec{s} is a leaf separation. Let $S_{\geq \vec{r}}$ be the set of separations $x \in S$ that have an orientation $\vec{x} \geq \vec{r}$. Since \vec{r} is a leaf separation in an irredundant S -tree over a set of stars we have by the previous comments that the image of α is contained in $\vec{S}_{\geq \vec{r}}$.

Given $x \in S_{\geq \vec{r}} \setminus \{r\}$ we have, since \vec{r} is non-trivial, that only one of the two orientations of x , say \vec{x} is such that $\vec{x} \geq \vec{r}$. So, we can define a function $f \downarrow_{\vec{s}}^{\vec{r}}$ on $\vec{S}_{\geq \vec{r}} \setminus \{\overleftarrow{r}\}$ by³

$$f \downarrow_{\vec{s}}^{\vec{r}}(\vec{x}) := \vec{x} \vee \vec{s} \text{ and } f \downarrow_{\vec{s}}^{\vec{r}}(\overleftarrow{x}) := (\vec{x} \vee \vec{s})^*.$$

Given an S -tree (T, α) and $\vec{s} \geq \vec{r}$ as above let $\alpha' := f \downarrow_{\vec{s}}^{\vec{r}} \circ \alpha$. The *shift of* (T, α) *onto* \vec{s} is the S -tree (T, α') .

We say that \vec{s} *emulates* \vec{r} in \vec{S} if $\vec{r} \leq \vec{s}$ and for every $\vec{t} \in \vec{S}_{\geq \vec{r}} \setminus \{\overleftarrow{r}\}$, $\vec{s} \vee \vec{t} \in \vec{S}$. Given a particular set of stars $\mathcal{F} \subseteq 2^{\vec{S}}$ we say further that \vec{s} *emulates* \vec{r} in \vec{S} for \mathcal{F} if \vec{s} emulates \vec{r} in \vec{S} and for any star $\sigma \subset \vec{S}_{\geq \vec{r}} \setminus \{\overleftarrow{r}\}$ in \mathcal{F} that contains an element $\vec{t} \geq \vec{r}$ we also have $f \downarrow_{\vec{s}}^{\vec{r}}(\sigma) \in \mathcal{F}$. The usefulness of this property is exhibited by the following lemma, which is key both in the proof of Theorem 5 from [7], and will be essential for the proof of our central lemma.

³The exclusion of \overleftarrow{r} here is for a technical reason, since it could be the case that $\vec{r} < \overleftarrow{r}$, however we want to insist that $f \downarrow_{\vec{s}}^{\vec{r}}(\overleftarrow{r})$ is the inverse of $f \downarrow_{\vec{s}}^{\vec{r}}(\vec{r})$




Figure 4: Shifting a separation $\vec{x} \geq \vec{r}$ under $f \downarrow_{\vec{s}}$.

Lemma 4. *[[7], Lemma 4.2] Let $(\vec{S}, \leq, *)$ be a separation system, $\mathcal{F} \subseteq 2^{\vec{S}}$ a set of stars, and let (T, α) be an irredundant S -tree over \mathcal{F} . Let \vec{r} be a nontrivial and nondegenerate separation which is a leaf separation of (T, α) , and is not the image of any other edge in T , and let \vec{s} emulate \vec{r} in \vec{S} . Then the shift of (T, α) onto \vec{s} is an S -tree over $\mathcal{F} \cup \{\{\vec{s}\}\}$ in which \vec{s} is a leaf separation, associated with a unique leaf.*

It is shown in [[7], Lemma 2.4] that if we have an S -tree over \mathcal{F} , (T, α) , and a set of non-trivial and non-degenerate leaf separations, \vec{r}_i , of (T, α) then there also exists an irredundant S -tree over \mathcal{F} , (T', α') , such that each \vec{r}_i is a leaf separation of (T', α') and is not the image of any other edge in T' .

We say a set $\mathcal{F} \subseteq 2^{\vec{S}}$ forces a separation \vec{r} if $\{\vec{r}\} \in \mathcal{F}$ or r is degenerate. Note that the non-degenerate forced separations in \mathcal{F} are precisely those separations which can appear as leaf separations in an S -tree over \mathcal{F} . We say \mathcal{F} is *standard* if it forces every trivial separation in \vec{S} .

We say that a separation system \vec{S} is *separable* if for any two non-trivial and non-degenerate separations $\vec{r}, \vec{r}' \in \vec{S}$ such that $\vec{r} \leq \vec{r}'$ there exists a separation $s \in S$ such that \vec{s} emulates \vec{r} in \vec{S} and \overleftarrow{s} emulates \vec{r}' in \vec{S} . We say that \vec{S} is \mathcal{F} -separable if for all non-trivial and non-degenerate $\vec{r}, \vec{r}' \in \vec{S}$ that are not forced by \mathcal{F} such that $\vec{r} \leq \vec{r}'$ there exists a separation $s \in S$ such that \vec{s} emulates \vec{r} in \vec{S} for \mathcal{F} and \overleftarrow{s} emulates \vec{r}' in \vec{S} for \mathcal{F} . Often one proves that \vec{S} is \mathcal{F} -separable in two steps, first by showing it is separable, and then by showing that \mathcal{F} is *closed under shifting*: that whenever \vec{s} emulates some \vec{r} in \vec{S} , it also emulates that \vec{r} in \vec{S} for \mathcal{F} .

We are now in a position to state the Strong Duality Theorem from [7].

Theorem 5. *[[7], Theorem 4.3] Let $(\vec{U}, \leq, *, \vee, \wedge)$ be a universe of separations containing a separation system $(\vec{S}, \leq, *)$. Let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a standard set of stars. If \vec{S} is \mathcal{F} -separable, exactly one of the following assertions holds:*

- *There exists an S -tree over \mathcal{F} .*
- *There exists an \mathcal{F} -tangle of S .*

The property of being \mathcal{F} -separable may seem a rather strong condition to hold, however in [8] it is shown that for all the sets \mathcal{F} describing classical ‘large’ objects (such as tangles or brambles) the separation systems \vec{S}_k are \mathcal{F} -separable. More specifically, by definition a k -tangle is a consistent orientation which avoids the set \mathcal{T}_k as defined earlier. In fact it is shown in [8] that a consistent orientation avoids \mathcal{T}_k if and only if it avoids the set of stars in \mathcal{T}_k

$$\mathcal{T}_k^* = \{ \{(A_i, B_i)\}_1^3 : \{(A_i, B_i)\}_1^3 \subseteq S_k \text{ is a star and } \bigcup_i G[A_i] = G \}.$$

Note that \mathcal{T}_k^* is standard. Indeed it forces all the small separations (A, V) , and so it forces the trivial separations. It can also be checked that \vec{S}_k is \mathcal{T}_k^* -separable.

The dual structure to a k -tangle is therefore an S_k -tree over \mathcal{T}_k^* . It is shown in [8] that the existence of such an S_k -tree is equivalent to the existence of a branch-decomposition of width $< k$ for all $k \geq 3$. We note that the condition that $k \geq 3$ is due to a quirk in how branch-width is traditionally defined, which results in, for example, stars having branch-width 1 but all other trees having branch-width 2, whilst both contain 2-tangles.

If a tree-decomposition (T, \mathcal{V}) of a graph G is such that the set of separations induced by the edges of T is an S_k -tree over \mathcal{T}_k^* for some k , then there is some smallest such k' , and we say the *branch-width* of the tree-decomposition is $k' - 1$. If no such k exists then we will let the branch-width be infinite. By the preceding discussion we have that the branch-width (in the traditional sense) of a graph is the smallest k such that G has a tree-decomposition of branch-width k (except when the branch-width of G is 1), and so this should not cause too much confusion.

2.3 Canonical tree-Decompositions distinguishing tangles

Given two orientations O_1 and O_2 of a set of separations S we say that a separation s *distinguishes* O_1 and O_2 if $\vec{s} \in O_1$ and $\overleftarrow{s} \in O_2$. As in the previous section, every tree-decomposition, (T, \mathcal{V}) , corresponds to some nested set of separations, \mathcal{N} . We say that a tree-decomposition *distinguishes* O_1 and O_2 if there is some separation in \mathcal{N} which distinguishes O_1 and O_2 . If O_1 and O_2 are consistent, then the tree-decomposition will distinguish them if and only if they are contained in different parts of the tree.

As in Section 2.2 a k -block b can be viewed as an orientation of S_k . Indeed given any separation (A, B) with $\text{ord}(A, B) < k$, since b is $(< k)$ -inseparable, $b \subseteq A$ or $b \subseteq B$, so we can think of b as orienting each $s \in S_k$ towards the side of the separations that b lies in. In [3] Carmesin, Diestel, Hundertmark and Stein showed how to algorithmically construct a nested set of separations in a graph G (and so a tree-decomposition) in a canonical way, that is, invariant with respect to the automorphism group of G , which distinguishes all of its k -blocks, for a given k .

These ideas were extended in [1] to construct canonical tree-decompositions which distinguish all the k -profiles in a graph, a common generalization of k -tangles and k -blocks. A k -*profile* can be defined as a \mathcal{P}_k -tangle of S_k , where

$$\mathcal{P}_k = \{\sigma = \{(A, B), (C, D), (B \cap D, A \cup C)\} : \sigma \subseteq \vec{S}_k\}.$$

More generally, given a universe of separations $(\vec{U}, \leq, *, \vee, \wedge)$ with an order function containing a separation system $(\vec{S}, \leq, *)$, we can define as before an S -*profile* to be a \mathcal{P}_S -tangle of S where

$$\mathcal{P}_S = \{\sigma = \{\vec{r}, \vec{s}, \overleftarrow{r} \wedge \overleftarrow{s}\} : \sigma \subseteq \vec{S}\}.$$

Given two distinct S -profiles P_1 and P_2 there is some $s \in S$ which distinguishes them. Furthermore, there is some minimal l such that there is a separation of order l which distinguishes P_1 and P_2 , and we define $\kappa(P_1, P_2) := l$. We say that a separation s distinguishes P_1 and P_2 *efficiently* if s distinguishes P_1 and P_2 and $|s| = \kappa(P_1, P_2)$. Given a set of profiles ϕ we say that a separation s is ϕ -*essential* if it efficiently distinguishes some pair of profiles in ϕ . We will often consider in particular, as in the case of graphs, the separation system arising from those separations in a universe of order $< k$, that is we define

$$\vec{S}_k = \{\vec{u} \in \vec{U} : |\vec{u}| < k\},$$

where in general it should be clear from the context which universe S_k lives in.

In [1] a number of different algorithms, which they call k -*strategies*, are described for constructing a nested set of separations distinguishing a set of profiles. These algorithms build the set of separations in a series of steps, and at each step there is a number of options for how to pick the next set of separations. A k -strategy is then a description of which choice to make at each step. The authors showed that, regardless of which choices are made at each step, this algorithm will produce a nested set of separations distinguishing all the profiles in G . We say a set of profiles is *canonical* if it is fixed under every automorphism of G . In particular the following is shown.

Theorem 6. *[[1] Theorem 4.4] Every k -strategy Σ determines for every canonical set ϕ of k -profiles of a graph G a canonical nested set $\mathcal{N}_\Sigma(G, \phi)$ of ϕ -essential separations of order $< k$ that distinguishes all the profiles in ϕ efficiently.*

Note that any k -tangle, O , is also a k -profile. Indeed, it is a simple check that O is consistent. Also for any pair of separations $(A, B), (C, D) \in \vec{S}_k$ we have that $G[A] \cup G[C] \cup G[B \cap D] = G$, since any edge not contained in A or C is contained in both B and D . Hence, $\{(A, B), (C, D), (B \cap D, A \cup C)\} \in \mathcal{T}_k$, and so $\mathcal{P}_k \subseteq \mathcal{T}_k$. Therefore any k -tangle, which by definition avoids \mathcal{T}_k , must also avoid \mathcal{P}_k , and so must be a k -profile. Similarly one can show that the orientations defined by k -blocks are consistent and \mathcal{P}_k avoiding, and so k -profiles. Even more, there is some family $\mathcal{B}_k \supseteq \mathcal{P}_k$ such that the orientations defined by k -blocks are \mathcal{B}_k -tangles, and if there is a \mathcal{B}_k -tangle of S_k then the graph G contains a unique k -block corresponding to this orientation.

One of the aims of [7, 8] had been to develop a duality theorem which would be applicable to k -profiles and k -blocks. The same authors showed in [6] that there is a more general duality theorem of a similar kind which applies in these cases, however the dual objects in this theorem correspond to a more general object than the classical notion of tree-decompositions.

Nevertheless, it was posed as an open question whether or not there was a duality theorem for k -profiles or k -blocks expressible within the framework of [7]. By Theorem 5 it would be sufficient to show that there is a standard set of stars \mathcal{F} such that the set of k -profiles or k -blocks coincides with the set of \mathcal{F} -tangles. Recently Diestel, Eberenz and Erde [4] showed that, if we insist the orientations satisfy a slightly stronger consistency condition, this will be the case. We say that an orientation O of a separation system S is *regular* if whenever we have r and s such that $\vec{r} \leq \vec{s}$, O does not contain both \vec{r} and \vec{s} . We note that a consistent orientation is regular if and only if it contains every small separation. A *regular \mathcal{F} -tangle* of S is then a regular \mathcal{F} -avoiding orientation of S , and a *regular S -profile* is a regular \mathcal{P}_S -tangle. For most natural examples of separation systems there will not be a difference between regular and irregular profiles. Indeed, in [4] it is shown that for $k \geq 3$ every k -profile of a graph is in fact a regular k -profile⁴.

We say a separation system is *submodular* if whenever $\vec{r}, \vec{s} \in \vec{S}$ either $\vec{r} \wedge \vec{s}$ or $\vec{r} \vee \vec{s} \in \vec{S}$. Note that, if a universe U has an order function, then the separation systems S_k are submodular.

Theorem 7. [Diestel, Eberenz and Erde [4]] *Let S be a separable submodular separation system contained in some universe of separations $(\vec{U}, \leq, *, \vee, \wedge)$, and let $\mathcal{F} \supseteq \mathcal{P}_S$. Then there exists a standard set of stars \mathcal{F}^* (which is closed under shifting, and contains $\{\vec{r}\}$ for every co-small \vec{r}) such that every regular \mathcal{F} -tangle of S is an \mathcal{F}^* -tangle of S , and vice versa, and such that the following are equivalent:*

- *There is no regular \mathcal{F} -tangle of S ;*
- *There is no \mathcal{F}^* -tangle of S ;*
- *There is an S -tree over \mathcal{F}^* .*

In the case where $S = S_k$ is the set of separations of a graph with $k \geq 3$, we have that $\mathcal{F} \supseteq \mathcal{P}_k$, and so every \mathcal{F} -tangle is a k -profile, and so regular. Hence, in this case, we can omit the word regular from the statement of the theorem. We note that $\mathcal{T}_k \supseteq \mathcal{P}_k$, (and in fact the \mathcal{T}_k^* of the theorem can be taken to be the \mathcal{T}_k^* defined earlier) and so Theorem 7 also implies the tangle/branch-width duality theorem.

Applying the result to \mathcal{P}_k or \mathcal{B}_k also gives a duality theorem for k -blocks and k -profiles. As in the case of tangles, if a tree-decomposition (T, \mathcal{V}) of a graph G is such that the set of separations induced by the edges of T is an S_k -tree over \mathcal{P}_k^* for some k , then there is some smallest such k' , and we say the *profile-width* of the tree-decomposition is $k' - 1$. If no such k exists then we will let the profile-width be infinite. The *profile-width* of a graph is then the smallest k such that G has a tree-decomposition of profile-width k . Then, as was the case with tangles, Theorem 7 tells us that the profile-width of a graph is the largest k such that G contains a k -profile. We define the *block-width* of a tree-decomposition and graph in the same way.

In a similar way as before, we can think of any part in a tree-decomposition of block-width at most $k - 1$ as being ‘too small’ to contain a k -block, as the corresponding star of separations must lie in \mathcal{B}_k^* , and by Theorem 7 every k -block defines an orientation of S_k which avoids \mathcal{B}_k^* .

⁴There do exist pathological examples of 2-profiles in graphs which are not regular, however they can be easily characterized.

3 Refining a tree-decomposition

Given a set of profiles of a graph, ϕ , we say a part V_t of a tree-decomposition is ϕ -essential if some profile from ϕ is contained in this part. We will keep in mind as a motivating example the case $\phi = \tau_k$, the set of k -tangles and, when the set of profiles considered is clear, we will refer to such parts simply as *essential*. Conversely if no such profile is contained in the part we call it *inessential*. The main result of the paper can now be stated formally.

Lemma 8. *Let $(\vec{U}, \leq, *, \vee, \wedge)$ be a universe of separations with an order function. Let ϕ be a set of S_k -profiles and let \mathcal{F} be a standard set of stars which contains $\{\vec{r}\}$ for every co-small \vec{r} , and which is closed under shifting, such that ϕ is the set of \mathcal{F} -tangles. Let $\sigma = \{\vec{s}_i : i \in [n]\} \subseteq \vec{S}_k$ be a non-empty star of separations such that each s_i is ϕ -essential, and let $\mathcal{F}' = \mathcal{F} \cup \bigcup_1^n \{\vec{s}_i\}$*

Then either there is an \mathcal{F}' -tangle of S_k , or there is an S_k -tree over \mathcal{F}' in which each \vec{s}_i appears as a leaf separation.

If we compare Lemma 8 to Theorem 5, we see that Lemma 8 can be viewed in some way as a method of building a new duality theorem from an old one, by adding some singleton separations to our set \mathcal{F} . The restriction to considering only S_k -profiles rather than those of arbitrary separation systems S contained in U comes from the proof, where we need to use the submodularity of the order function to show that certain separations emulate others. It would be interesting to know if the result would still be true for any S which is separable, or even any pair \mathcal{F} and S such that S is \mathcal{F} -separable. The condition that \mathcal{F} contains every co-small separation as a singleton is to ensure that the \mathcal{F} -tangles are regular \mathcal{F} -tangles, as we will need to use the slightly stronger consistency condition in the proof.

What does Lemma 8 say in the case of k -tangles arising from graphs? Recall that τ_k is the set of \mathcal{T}_k^* -tangles, and that \mathcal{T}_k^* is closed under shifting, and contains $\{\vec{r}\}$ for every co-small \vec{r} . Given a star $\sigma = \{\vec{s}_i : i \in [n]\} \subseteq \vec{S}_k$ we note that a $\mathcal{T}_k^* \cup \bigcup_1^n \{\vec{s}_i\}$ -tangle is just a \mathcal{T}_k^* -tangle which contains \vec{s}_i for each i , and so it is a k -tangle which orients the star inwards. Conversely, an S_k -tree over $\mathcal{T}_k^* \cup \bigcup_1^n \{\vec{s}_i\}$ in which each \vec{s}_i appears as a leaf separation will give a tree-decomposition of the part of the graph at σ . In particular, since each of the separations in the tree will be nested with σ , the separators $A_i \cap B_i$ of the separations \vec{s}_i will lie entirely on one side of every separation in the tree, and so this will in fact be a decomposition of the torso of the part (since any extra edges in the torso lie inside the separators).

Therefore, in practice this tells us that if we have a part in a tree-decomposition whose separators are τ_k -essential then either there is a k -tangle in the graph which is contained in that part, or there is a tree-decomposition of the torso of that part with branch-width $< k$. In the second case we can then refine the original tree-decomposition by combining it with this new tree-decomposition. By applying this to each inessential part of one of the canonical tree-decompositions formed in [1] we get the following result, which easily implies Theorems 2 and 3 by taking $\mathcal{F} = \mathcal{T}_k^*$ and \mathcal{B}_k^* respectively.

Corollary 9. *Let $k \geq 3$ and let $\mathcal{F} \supseteq \mathcal{P}_k$ be such that the set ϕ of regular \mathcal{F} -tangles is canonical. If \mathcal{F}^* is defined as in Theorem 7 then there exists a nested set of separations $\mathcal{N} \subseteq S_k$ corresponding to an S_k -tree (T, α) of G such that:*

- *there is a subset $\mathcal{N}' \subseteq \mathcal{N}$ that is fixed under every automorphism of G and distinguishes all the regular \mathcal{F} -tangles in ϕ efficiently;*
- *every vertex $t \in T$ either contains a regular \mathcal{F} -tangle or $\{\alpha(t', t) : (t', t) \in \vec{ET}\} \in \mathcal{F}^*$.*

Proof. By Theorem 6 there exists a canonical nested set \mathcal{N}' of ϕ -essential separations of order $< k$ that distinguishes all the regular \mathcal{F} -tangles in ϕ efficiently, and by Theorem 7 ϕ is also the set of \mathcal{F}^* -tangles. Given an inessential part V_t in the corresponding tree-decomposition (T, \mathcal{V}) , this part corresponds to some star of separations $\sigma = \{\vec{s}_i : i \in [n]\} \subseteq \mathcal{N}'$. Each $\vec{s}_i \in \mathcal{N}'$ is ϕ -essential, and, by Theorem 7, \mathcal{F}^* is a standard set of stars which is closed under shifting, and contains $\{\vec{r}\}$ for every co-small \vec{r} . Hence, by Lemma 8, if we let $\mathcal{F}' = \mathcal{F}^* \cup \bigcup_1^n \{\vec{s}_i\}$, there is either an \mathcal{F}' -tangle of S_k , or an S_k -tree over \mathcal{F}' in which each \vec{s}_i appears as a leaf separation.

Suppose that there exists an \mathcal{F}' -tangle O . Since O avoids $\mathcal{F} \supseteq \mathcal{F}^*$, it is also an \mathcal{F}^* -tangle, and so $O \in \phi$. By assumption \mathcal{N}' distinguishes all the regular \mathcal{F} -tangles in ϕ , so O is contained in some part of the tree-decomposition, and since O avoids $\{\{\vec{s}_i\} : i \in [n]\}$, it must extend σ , and so this part must be V_t . However, this contradicts the assumption that V_t is inessential.

Therefore, by Lemma 8, there exists an S_k -tree over $\mathcal{F}^* \cup \bigcup_1^n \{\overleftarrow{s_i}\}$. This gives a nested set of separations \mathcal{N}_t which contains the set σ . If we take such a set for each inessential V_t then the set

$$\mathcal{N} = \mathcal{N}' \cup \bigcup_{V_t \text{ inessential}} \mathcal{N}_t$$

satisfies the conditions of the corollary. \square

We note that, whilst the existence of such a tree-decomposition is interesting in its own right, perhaps a more useful application of Lemma 8 is that we can conclude the same for *every* tree-decomposition constructed by the algorithms in [1]. So, we are able to choose whichever algorithm we want to construct our initial tree-decomposition, perhaps in order to have some control over the structure of the essential parts, and we can still conclude that the inessential parts have small branch-width.

Apart from the set τ_k of k -tangles there is another natural set of tangles for which tangle-distinguishing tree-decompositions have been considered. Since a k -tangle, as a \mathcal{T}_k -avoiding orientation of S_k , induces an orientation on S_i for all $i \leq k$, it induces an i -tangle for all $i \leq k$. If an i -tangle for some i is not induced by any k -tangle with $k > i$ we say it is a *maximal tangle*.

Robertson and Seymour [9] showed that there is a decomposition of the graph which distinguishes its maximal tangles, but the theorem does not tell us much about the structure of this tree-decomposition. The approach of Carmesin et al was extended by Diestel, Hundertmark and Lemanczyk [5] to show how an iterative approach to Theorem 6 could be used to build canonical tree-decompositions distinguishing the maximal tangles in a graph (in fact they showed a stronger result for a broader class of profiles which implies the result for tangles). In particular, the results of [5] imply the following.

Theorem 10. *If ϕ is a canonical set of tangles in a graph G , then there exists a canonical nested set $\mathcal{N}(G, \phi)$ of ϕ -essential separations that distinguishes all the tangles in ϕ efficiently.*

In particular we can apply this to the set of maximal tangles. By looking directly at the proof in [5] one can see the structure of the tree-decomposition formed. The proof proceeds iteratively, by choosing for each i in a turn a nested set of $(i - 1)$ -separations (that is, separations of order $(i - 1)$), which distinguishes efficiently the pairs of i -tangles which are distinguished efficiently by an $(i - 1)$ -separation, such that this set is also nested with the previously constructed sets.

At each stage in the construction we have a tree-decomposition which distinguishes all the tangles of order $\leq i$ in the graph. Some of these i -tangles however will extend to $(i + 1)$ -tangles in different ways (induced by distinct maximal tangles in the graph). The next stage constructs a nested set of separations distinguishing such tangles, which gives a tree-decomposition of the torsos of the relevant parts. In these tree-decompositions some parts will be ‘essential’, and containing $(i + 1)$ -tangles, but some will be inessential.

It is natural to expect that the inessential parts constructed at stage i should have branch-width $< i$, by a similar argument as Corollary 9. However it is not always the case that the separators of the inessential part satisfy the conditions of Lemma 8, since it can be the case that these inessential parts have separators which are separations constructed in an earlier stage of the process, and as such might not efficiently distinguish a pair of tangles of order i .

Question 11. *Can we bound the branch-width of the inessential parts in such a tree-decomposition in a similar way?*

A positive answer to the previous question in the strongest form would give the following analogue of Theorem 2.

Conjecture 12. *For every graph G there exists a canonical sequence of tree-decompositions (T_i, \mathcal{V}_i) for $1 \leq i \leq n$ of G such that*

- (T_i, \mathcal{V}_i) distinguishes every i -tangle in G for each i ;
- (T_n, \mathcal{V}_n) distinguishes the set of maximal tangles in G .
- $(T_{i+1}, \mathcal{V}_{i+1})$ refines (T_i, \mathcal{V}_i) for each i ;
- The torso of every inessential part in (T_i, \mathcal{V}_i) has branch-width $< i$.

3.1 Proof of Lemma 8

Proof of Lemma 8. Let us write

$$\overline{\mathcal{F}} = \mathcal{F} \cup \{\{\overleftarrow{x}\} : \overleftarrow{s}_i \leq \overleftarrow{x} \text{ for some } i \in [n]\}.$$

We first claim that \overrightarrow{S}_k is $\overline{\mathcal{F}}$ -separable. We note that by [[8], Lemma 3.4] for every universe \overrightarrow{U} and any $k \in \mathbb{N}$, the separation system \overrightarrow{S}_k is separable. Therefore it is sufficient to show that $\overline{\mathcal{F}}$ is closed under shifting. By assumption \mathcal{F} is closed under shifting, and the image of any singleton star $\{\overleftarrow{x}\} \in \overline{\mathcal{F}}$ under some relevant $f \downarrow_{\overrightarrow{S}}^{\overrightarrow{U}}$ is $\{\overleftarrow{y}\}$ for some separation $\overleftarrow{x} \leq \overleftarrow{y}$, and hence $\{\overleftarrow{y}\} \in \overline{\mathcal{F}}$. Therefore, $\overline{\mathcal{F}}$ is closed under shifting. Furthermore, since \mathcal{F} was standard, so is $\overline{\mathcal{F}}$. Hence, we can apply Theorem 5 to $\overline{\mathcal{F}}$.

By Theorem 5, either there exists an S_k -tree over $\overline{\mathcal{F}}$, or there exists an $\overline{\mathcal{F}}$ -tangle. Since $\overline{\mathcal{F}} \supset \mathcal{F}'$, every $\overline{\mathcal{F}}$ -tangle is also an \mathcal{F}' -tangle, and so in the second case we are done. Therefore we may assume that there exists an S_k -tree over $\overline{\mathcal{F}}$, (T, α) . We will use (T, α) to form an S_k -tree over \mathcal{F}' .

Since there is no $\overline{\mathcal{F}}$ -tangle, each \mathcal{F} -tangle O must contain some \overleftarrow{s}_i . We note that, since by assumption \mathcal{F} contains every co-small separation, O is regular. Hence, since σ is a star, this \overleftarrow{s}_i is unique. We claim that, for every \mathcal{F} -tangle O such that $\overleftarrow{s}_i \in O$ there is some leaf separation $\overleftarrow{x} \in \alpha(\overrightarrow{E}(T))$ such that $\overleftarrow{x} \leq \overleftarrow{s}_i$.

Indeed, since O is a consistent orientation of \overrightarrow{S}_k , it is contained in some vertex of (T, α) . However, the star of separations at that vertex, by definition of an \mathcal{F} -tangle, cannot lie in \mathcal{F} , and so must lie in $\overline{\mathcal{F}} \setminus \mathcal{F}$. Since each of these stars are singletons, the vertex must be a leaf. Therefore, there is some leaf separation \overleftarrow{x} such that $\overleftarrow{x} \in O$. Since $\{\overleftarrow{x}\} \in \overline{\mathcal{F}} \setminus \mathcal{F}$, it follows that $\overleftarrow{s}_r \leq \overleftarrow{x}$ for some $r \in [n]$. However, since $\overleftarrow{s}_i \in O$, and it was the unique separation in σ with that property, it follows that $r = i$, and so $\overleftarrow{s}_i \leq \overleftarrow{x}$ as claimed.

If the only leaf separations in $\overline{\mathcal{F}} \setminus \mathcal{F}$ were the separations $\{\overleftarrow{s}_i : i \in [n]\}$ then (T, α) would be the required S -tree over \mathcal{F}' . In general however the tree will have a more arbitrary set $\{\overleftarrow{x}_{i,j}\}$ of leaf separations (along with some leaf separations arising as separations forced by \mathcal{F}) where $\overleftarrow{s}_i \leq \overleftarrow{x}_{i,j}$, see Figure 5. Note that there may not necessarily be any edges in this tree corresponding to the separations s_i .




Figure 5: The S_k -tree over $\overline{\mathcal{F}}$ with unlabelled leaves corresponding to separations forced by \mathcal{F} .

We claim that each \overrightarrow{s}_i emulates some $\overrightarrow{x}_{i,j}$ in \overrightarrow{S}_k for $\overline{\mathcal{F}}$. By assumption, every $s_i \in \sigma$ distinguishes efficiently some pair O_1 and O_2 of \mathcal{F} -tangles. Suppose that $\overrightarrow{s}_i \in O_1$ and $\overleftarrow{s}_i \in O_2$. By our previous claim, there is some leaf separation $\overrightarrow{x}_{i,j}$ such that $\overleftarrow{x}_{i,j} \in O_2$. We claim that \overrightarrow{s}_i emulates this $\overrightarrow{x}_{i,j}$ in \overrightarrow{S}_k . Note that, since $\overline{\mathcal{F}}$ is closed under shifting, it would follow that \overrightarrow{s}_i emulates $\overrightarrow{x}_{i,j}$ in \overrightarrow{S}_k for $\overline{\mathcal{F}}$. Note that, since \overrightarrow{s}_i and $\overrightarrow{x}_{i,j}$ both distinguish two \mathcal{F} -tangles, they are non-trivial and non-degenerate.

Indeed, given any separation $\vec{r} \geq \vec{x}_{i,j}$ we have that $\vec{s}_i \geq \vec{s}_i \wedge \vec{r} \geq \vec{x}_{i,j}$ and so $\vec{s}_i \wedge \vec{r}$ distinguishes O_1 and O_2 . Therefore, since s_i distinguishes O_1 and O_2 efficiently, $|\vec{s}_i \wedge \vec{r}| \geq |\vec{s}_i|$. Hence, by submodularity, $|\vec{s}_i \vee \vec{r}| \leq |\vec{r}| < k$ and so $\vec{s}_i \vee \vec{r} \in S_k$. Therefore the image of $f \downarrow_{\vec{s}_i}^{\vec{x}_{i,j}}$ is contained in S_k and so \vec{s}_i emulates $\vec{x}_{i,j}$ in \vec{S}_k . Furthermore, since $\vec{x}_{i,j}$ is non-trivial and non-degenerate, by the comment after Lemma 4 we can assume that T is irredundant, and that $\vec{x}_{i,j}$ is not the image of any other edge in T .

Since \vec{s}_n and $\vec{x}_{n,j}$ satisfy the conditions of Lemma 4, we conclude that the shift of (T, α) onto \vec{s}_n is an S_k -tree over $\vec{\mathcal{F}}$ which contains \vec{s}_n as a leaf separation, and not as the image of any other edge. Let us write (T_n, α_n) for this S_k -tree.

If there is some leaf separation \overleftarrow{r} of (T_n, α_n) such that $\overleftarrow{s}_n < \overleftarrow{r}$ then, since the leaf separations form a star and \overleftarrow{s}_i is the image of a unique leaf, we also have that $\overleftarrow{s}_i < \overleftarrow{r}$. Hence, \vec{r} is trivial, and so $\{\overleftarrow{r}\} \in \mathcal{F}$. Therefore (T_n, α_n) is also an S_k -tree over

$$\vec{\mathcal{F}}_n = \mathcal{F} \cup \{\overleftarrow{s}_n\} \cup \{\{\overleftarrow{x}\} : \overleftarrow{s}_i \leq \overleftarrow{x} \text{ for some } i \in [n-1]\}.$$

If we repeat this argument for each $1 \leq i \leq n$, we end up with a sequence of S_k -trees $(T_n, \alpha_n), (T_{n-1}, \alpha_{n-1}), \dots, (T_1, \alpha_1)$ over $\vec{\mathcal{F}}$ such that (T_j, α_j) is also an S_k -tree over

$$\vec{\mathcal{F}}_j = \mathcal{F} \cup \{\overleftarrow{s}_i : i \geq j\} \cup \{\{\overleftarrow{x}\} : \overleftarrow{s}_i \leq \overleftarrow{x} \text{ for some } i \in [j-1]\}.$$

We note that $\vec{\mathcal{F}}_1 = \mathcal{F}'$, and so (T_1, α_1) is an S_k -tree over \mathcal{F}' , completing the proof. \square

4 Further refining essential parts of tangle-distinguishing tree-decompositions

In some sense the tree-decompositions of Corollary 9 tell us most about the structure of the graph when the essential parts correspond closely to the profiles inside them. However, as the example in Figure 2 shows, sometimes there can be essential parts which could be further refined, in order to more precisely exhibit the structure of the graph.

In this section we will discuss how the tools from the paper can be used to achieve this goal. Given a graph G we call a separation $\overleftarrow{x} \in \vec{S}_k$ *inessential* if $\overleftarrow{x} \in O$ for every k -tangle O of G . Given a k -tangle O let $\mathcal{M}(O)$ be the set of maximal separations in O , and let $\mathcal{M}_I(O)$ be the set of maximal inessential separations. Our main tool will be the following lemma.

Lemma 13. *Let G be a graph, O be a k -tangle of G and let $\overleftarrow{x} \in \mathcal{M}_I(O)$ be non-trivial. Then there is an S_k -tree over $\mathcal{T}_k^* \cup \{\overleftarrow{x}\}$.*

Proof. As in the proof of Lemma 8 let us consider the family of stars

$$\mathcal{F} = \mathcal{T}_k^* \cup \{\{\overleftarrow{r}\} : \overleftarrow{x} \leq \overleftarrow{r}\}.$$

A similar argument show that this family is standard and closed under shifting, and so Theorem 5 asserts the existence of a \mathcal{F} -tangle, or an S_k -tree over \mathcal{F} . As before, an \mathcal{F} -tangle would be a k -tangle of G which contains \overleftarrow{x} , contradicting the fact that \overleftarrow{x} is inessential. Therefore, there is an S_k -tree over \mathcal{F} . However, O must live in some part of this tree-decomposition, and since O is \mathcal{T}_k^* -avoiding it must live in some leaf vertex, corresponding to a singleton star $\{\overleftarrow{r}\}$ for some $\overleftarrow{x} \leq \overleftarrow{r}$. However, \overleftarrow{x} was a maximal separation in O and hence $\overleftarrow{r} \notin O$ unless $\overleftarrow{r} = \overleftarrow{x}$. Therefore the S_k -tree is in fact over $\mathcal{T}_k^* \cup \{\overleftarrow{x}\}$. \square

Lemma 13 tell us that for every $\overleftarrow{x} \in \mathcal{M}_I(O)$ there is a tree-decomposition of the part of the graph behind \overleftarrow{x} with branch-width $< k$. So, we could perhaps hope to refine our canonical k -tangle-distinguishing tree-decompositions further using these tree-decompositions. However, there is no guarantee that $\mathcal{M}_I(O)$ will be nested with the τ_k -essential separations used in a k -tangle-distinguishing tree-decomposition, and so we cannot in general refine such tree-decompositions naively in this way. Moreso, in order to decompose as much of the inessential parts of the graph as possible we would like to take such a tree for *each* such maximal separation, however again in general, $\mathcal{M}_I(O)$ itself may not be nested.

Our plan will be to find, for each $\overleftarrow{x} \in \mathcal{M}_I(O)$, some inessential separation \overleftarrow{u} such that \overrightarrow{u} emulates \overrightarrow{x} in \overrightarrow{S}_k , that is also nested with the separations from the k -tangle-distinguishing tree-decomposition. Furthermore we would like to be able to do this in such a way that the separations \overrightarrow{u} emulating different maximal separations form a star. Then, for each maximal separation, we could shift the S_k -tree given by Lemma 13 to an S_k -tree over $\mathcal{T}_k^* \cup \{\overleftarrow{u}\}$. These tree-decompositions could then be used to refine our k -tangle-distinguishing tree-decomposition further. We will in fact show a more general result that may be of interest in its own right.

4.1 Uncrossing sets of separations

Given two separations $\overrightarrow{r} \leq \overrightarrow{s}$ in an arbitrary universe with an order function, we say that \overrightarrow{s} is *linked* to \overrightarrow{r} if for every $\overrightarrow{x} \geq \overrightarrow{r}$ we have that

$$|\overrightarrow{x} \vee \overrightarrow{s}| \leq |\overrightarrow{x}|.$$

In particular we note that if $\overrightarrow{r}, \overrightarrow{s} \in \overrightarrow{S}_k$, then \overrightarrow{s} being linked to \overrightarrow{r} implies that \overrightarrow{s} emulates \overrightarrow{r} in \overrightarrow{S}_k . We first note explicitly a fact used in the proof of Lemma 8.

Lemma 14. *Let $(\overrightarrow{U}, \leq, *, \vee, \wedge)$ be a universe of separations with an order function, and let $\overleftarrow{s} \leq \overleftarrow{r}$ be two separations in \overrightarrow{U} . If \overleftarrow{x} is a separation of minimal order such that $\overleftarrow{s} \leq \overleftarrow{x} \leq \overleftarrow{r}$, then \overrightarrow{x} is linked to \overrightarrow{r} .*

Proof. Given any separation $\overrightarrow{y} > \overrightarrow{r}$ we note that

$$\overleftarrow{s} \leq \overleftarrow{x} \leq \overleftarrow{x} \vee \overleftarrow{y} \leq \overleftarrow{r},$$

and so by minimality of \overleftarrow{x} we have that $|\overleftarrow{y} \vee \overleftarrow{x}| \geq |\overleftarrow{x}|$. Hence, by submodularity $|\overrightarrow{y} \vee \overrightarrow{x}| = |\overleftarrow{y} \wedge \overleftarrow{x}| \leq |\overrightarrow{y}|$, and so \overrightarrow{x} is linked to \overrightarrow{r} . Note that, by symmetry, \overleftarrow{x} is linked to \overleftarrow{s} also. \square

In what follows we will need to use two facts about a universe of separations. The first is true for any universe of separations, that for any two separations \overleftarrow{x} and \overleftarrow{y}

$$(\overleftarrow{x} \wedge \overleftarrow{y})^* = \overrightarrow{x} \vee \overrightarrow{y} \text{ and } (\overleftarrow{x} \vee \overleftarrow{y})^* = \overrightarrow{x} \wedge \overrightarrow{y}.$$

The second will not be true in general, and so we say a universe of separations is *distributive* if for every three separations $\overleftarrow{x}, \overleftarrow{y}$ and \overleftarrow{z} it is true that

$$(\overleftarrow{x} \wedge \overleftarrow{y}) \vee \overleftarrow{z} = (\overleftarrow{x} \vee \overleftarrow{z}) \wedge (\overleftarrow{y} \vee \overleftarrow{z}) \text{ and } (\overleftarrow{x} \vee \overleftarrow{y}) \wedge \overleftarrow{z} = (\overleftarrow{x} \wedge \overleftarrow{z}) \vee (\overleftarrow{y} \wedge \overleftarrow{z}).$$

It is a simple check that the universe of separations of a graph is distributive.

Lemma 15. *Let $(\overrightarrow{U}, \leq, *, \vee, \wedge)$ be a distributive universe of separations with an order function, and let \overleftarrow{x}_1 and \overleftarrow{x}_2 be two separations in \overrightarrow{U} . Let \overleftarrow{u}_1 be any separation of minimal order such that $\overleftarrow{x}_1 \wedge \overleftarrow{x}_2 \leq \overleftarrow{u}_1 \leq \overleftarrow{x}_1$ and let $\overleftarrow{u}_2 = \overleftarrow{x}_2 \wedge \overleftarrow{u}_1$. Then the following statements hold:*

- \overrightarrow{u}_1 is linked to \overrightarrow{x}_1 and \overrightarrow{u}_2 is linked to \overrightarrow{x}_2 ;
- $|\overleftarrow{u}_1| \leq |\overleftarrow{x}_1|$ and $|\overleftarrow{u}_2| \leq |\overleftarrow{x}_2|$;
- $\overleftarrow{u}_1 = \overleftarrow{x}_1 \wedge \overrightarrow{u}_2$.

Proof. We note that \overrightarrow{u}_1 is linked to \overrightarrow{x}_1 by Lemma 14. We want to show that \overrightarrow{u}_2 is linked to \overrightarrow{x}_2 , that is, given any $\overrightarrow{r} > \overrightarrow{x}_2$ we need that $|\overrightarrow{r} \vee \overrightarrow{u}_2| \leq |\overrightarrow{r}|$. We first claim that $\overrightarrow{r} \vee \overrightarrow{u}_2 = \overleftarrow{u}_1 \vee \overrightarrow{r}$. Indeed,

$$\overrightarrow{r} \vee \overrightarrow{u}_2 = \overrightarrow{r} \vee (\overrightarrow{x}_2 \vee \overleftarrow{u}_1) = (\overrightarrow{r} \vee \overrightarrow{x}_2) \vee \overleftarrow{u}_1 = \overrightarrow{r} \vee \overleftarrow{u}_1.$$

We also claim that $\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leq \overleftarrow{u}_1 \wedge \overrightarrow{r} \leq \overleftarrow{x}_1$. Indeed, $\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leq \overleftarrow{u}_1$ and $\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leq \overrightarrow{x}_2 \leq \overrightarrow{r}$ and so

$$\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leq \overleftarrow{u}_1 \wedge \overrightarrow{r} \leq \overleftarrow{u}_1 \leq \overleftarrow{x}_1.$$

Therefore, by minimality of \overleftarrow{u}_1 we have that $|\overleftarrow{u}_1 \wedge \overrightarrow{r}| \geq |\overleftarrow{u}_1|$ and so, by submodularity, it follows that

$$|\overrightarrow{r} \vee \overrightarrow{u}_2| = |\overleftarrow{u}_1 \vee \overrightarrow{r}| \leq |\overrightarrow{r}|,$$

as claimed.

By minimality of \overleftarrow{u}_1 we have that $|\overleftarrow{u}_1| \leq |\overleftarrow{x}_1|$. Also we note that, since $\overleftarrow{u}_2 = \overleftarrow{x}_2 \wedge \overrightarrow{u}_1$ we have that

$$|\overleftarrow{u}_2| + |\overleftarrow{x}_2 \vee \overrightarrow{u}_1| \leq |\overleftarrow{x}_2| + |\overleftarrow{u}_1|.$$

However, $|\overleftarrow{x}_2 \vee \overrightarrow{u}_1| = |\overleftarrow{u}_1 \wedge \overrightarrow{x}_2|$, and we claim that

$$\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leq \overleftarrow{u}_1 \wedge \overrightarrow{x}_2 \leq \overleftarrow{x}_1.$$

Indeed, that second inequality is clear since, $\overleftarrow{u}_1 \leq \overleftarrow{x}_1$. For the first we note that $\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leq \overrightarrow{x}_2$, and also $\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leq \overleftarrow{x}_1 \wedge \overrightarrow{u}_1 = \overleftarrow{u}_1$, and so $\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leq \overleftarrow{u}_1 \wedge \overrightarrow{x}_2$. Hence, by the minimality of \overleftarrow{u}_1 , we have $|\overleftarrow{x}_2 \vee \overrightarrow{u}_1| \geq |\overleftarrow{u}_1|$. Hence it follows that $|\overleftarrow{u}_2| \leq |\overleftarrow{x}_2|$, as claimed.

For the last condition, we have that $\overleftarrow{u}_1 \leq \overleftarrow{x}_1$ and $\overleftarrow{u}_1 \leq \overleftarrow{u}_1 \vee \overrightarrow{x}_2 = \overrightarrow{u}_2$, and so $\overleftarrow{u}_1 \leq \overleftarrow{x}_1 \wedge \overrightarrow{u}_2$. However,

$$\overleftarrow{x}_1 \wedge \overrightarrow{u}_2 = \overleftarrow{x}_1 \wedge (\overrightarrow{x}_2 \vee \overleftarrow{u}_1) = (\overleftarrow{x}_1 \wedge \overrightarrow{x}_2) \vee (\overleftarrow{x}_1 \wedge \overleftarrow{u}_1) = (\overleftarrow{x}_1 \wedge \overrightarrow{x}_2) \vee \overleftarrow{u}_1 \leq \overleftarrow{u}_1,$$

and so $\overleftarrow{u}_1 = \overleftarrow{x}_1 \wedge \overrightarrow{u}_2$. \square

We note that if we apply the above lemma to a pair of separations \overleftarrow{x}_1 and \overleftarrow{x}_2 such that x_1 distinguishes efficiently a pair of regular k -profiles, which x_2 does not distinguish, say $\overleftarrow{x}_1 \in P_1$ and $\overrightarrow{x}_1 \in P_2$ and $\overleftarrow{x}_2 \in P_1 \cap P_2$, then \overleftarrow{x}_1 is of minimal order over all separations $\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leq \overleftarrow{u}_1 \leq \overleftarrow{x}_1$. Hence, in Lemma 15, we can take $\overleftarrow{u}_1 = \overleftarrow{x}_1$ and $\overleftarrow{u}_2 = \overleftarrow{x}_2 \wedge \overrightarrow{x}_1$.

Indeed, suppose $\overleftarrow{x}_1 \wedge \overrightarrow{x}_2 \leq \overleftarrow{u}_1 \leq \overleftarrow{x}_1$ is of minimal order. We note that $\overleftarrow{u}_1 \in P_1$ by regularity. Similarly, let $\overleftarrow{u}_2 = \overleftarrow{x}_2 \wedge \overrightarrow{u}_1$ then $\overleftarrow{u}_2 \in P_2$ by regularity. Recall that, by Lemma 15 $\overleftarrow{u}_1 = \overleftarrow{x}_1 \wedge \overrightarrow{u}_2$. Hence, $\overrightarrow{u}_1 = \overrightarrow{x}_1 \vee \overleftarrow{u}_2 \in P_2$. Therefore, u_1 distinguishes P_1 and P_2 and so, by the efficiency of x_1 , $|\overleftarrow{x}_1| \leq |\overleftarrow{u}_1|$ as claimed.

The question remains as to what happens for a larger set of separations. It would be tempting to conjecture that the following extension of Lemma 15 holds, where we note that, in general, $(\overrightarrow{x} \wedge \overrightarrow{y}) \wedge \overrightarrow{z} = \overrightarrow{x} \wedge (\overrightarrow{y} \wedge \overrightarrow{z})$ and so, when writing such an expression we can, without confusion, omit the brackets.

Conjecture 16. *Let $(\overrightarrow{U}, \leq, *, \vee, \wedge)$ be a distributive universe of separations with an order function, and let $\{\overleftarrow{x}_i : i \in [n]\}$ be a set of separations in \overrightarrow{U} . Then there exists a set of separations $\{\overleftarrow{u}_i : i \in [n]\}$ such that the following conditions hold:*

- $\{\overleftarrow{u}_i : i \in [n]\}$ is a star;
- \overrightarrow{u}_i is linked to \overrightarrow{x}_i for all $i \in [n]$;
- $|\overleftarrow{u}_i| \leq |\overleftarrow{x}_i|$ for all $i \in [n]$;
- $\overleftarrow{u}_i = \overleftarrow{x}_i \wedge_{j \neq i} \overrightarrow{u}_j$ for all $i \in [n]$.

However, it seems difficult to ensure that the fourth condition holds with an inductive argument. We were able to show the following in the case of graph separations, by repeatedly applying Lemma 15. The extra sets $\{\overleftarrow{r}_i\}$ and ϕ appearing in the statement will be useful for the specific application we have in mind, the conclusion when these are empty is the weakened form of the above conjecture.

Lemma 17. *Let G be a graph, $k \geq 3$, and let ϕ be the set of k -profiles in G . Suppose that $\{\overleftarrow{r}_i = (A_i, B_i) : i \in [n]\}$ is a star composed of ϕ -essential separations, which distinguish efficiently some set ϕ' of regular k -profiles and let $\{\overleftarrow{x}_j = (X_j, Y_j) : j \in [m]\} \subseteq \overrightarrow{S}_k$ be such that $\overleftarrow{x}_j \in P$ for all $j \in [m]$ and $P \in \phi'$. Then there exists a set $\{\overleftarrow{u}_j : j \in [m]\}$ such that the following conditions hold:*

- $\{\overleftarrow{r}_i : i \in [n]\} \cup \{\overleftarrow{u}_j : j \in [m]\}$ is a star;
- $|\overleftarrow{u}_j| \leq |\overleftarrow{x}_j|$ for all $j \in [m]$;

- \vec{u}_j is linked to \vec{x}_j for all $j \in [m]$;
- $\vec{x}_j \wedge_i \vec{r}_i \wedge_{k \neq j} \vec{x}_k \leq \vec{u}_j \leq \vec{x}_j$ for all $j \in [m]$;
- $\bigcup_{j=1}^m X_j \cup \bigcup_{i=1}^n A_i = \bigcup_{j=1}^m U_j \cup \bigcup_{i=1}^n A_i$.

Proof. Let us start with a set of separations

$$Y = \{\overleftarrow{y}_i : i \in [n+m]\},$$

and some arbitrary order on the set of pairs $Y^{(2)}$. Initially we set $\overleftarrow{y}_i = \overleftarrow{x}_j$ for $j \in [m]$ and $\overleftarrow{y}_{m+i} = \overleftarrow{r}_i$ for $i \in [n]$. For each pair $\{\overleftarrow{y}_i, \overleftarrow{y}_j\}$ in order we apply Lemma 15 to this pair of separations and replace $\{\overleftarrow{y}_i, \overleftarrow{y}_j\}$ with the nested pair given by Lemma 15. After we have done this for each pair, we let $\overleftarrow{u}_j := \overleftarrow{y}_j$ for each $j \in [m]$.

Note that, since each \overleftarrow{x}_j is ϕ -inessential, and with each application of Lemma 15 we only ever replace a separation by one less than or equal to it, \overleftarrow{y}_j is also ϕ -inessential at each stage of this process for $j \in [m]$. Also, $\{\overleftarrow{r}_i : i \in [n]\}$ is a star, and so if we apply Lemma 15 to a pair \overleftarrow{r}_i and \overleftarrow{r}_k , neither is changed. Therefore, by the comment after Lemma 15, we may assume that at every stage in the process $\overleftarrow{y}_{m+i} = \overleftarrow{r}_i$ for each $i \in [n]$. In particular at the end of the process we have that

$$Y = \{\overleftarrow{r}_i : i \in [n]\} \cup \{\overleftarrow{u}_j : j \in [m]\}.$$

To see that the first condition is satisfied we note that, given any pair of separations \overleftarrow{y}_i and $\overleftarrow{y}_j \in Y$, at some stage in the process we applied Lemma 15 to this pair, and immediately after this step we have that $\overleftarrow{y}_i \leq \overleftarrow{y}_j$. Since Lemma 15 only ever replaces a separation with one less than or equal to it, it follows that at the end of the process Y is a star. Therefore the family $\{\overleftarrow{r}_i : i \in [n]\} \cup \{\overleftarrow{u}_j : j \in [m]\}$ forms a star.

To see that the second condition is satisfied we note that, whenever we apply Lemma 15 we only ever replace a separation with one whose order is less than or equal to the order of the original separation.

To see that the third condition is satisfied we note that whenever we apply Lemma 15 we only ever replace a separation with one whose inverse is linked to the inverse of the original separation. Therefore it would be sufficient to show that the property of being linked to is transitive. Indeed, suppose that $\vec{r} > \vec{s} > \vec{t}$, \vec{r} is linked to \vec{s} and \vec{s} is linked to \vec{t} , all in some separation system S . Let $\vec{x} > \vec{t}$ also be in S_k .

However, since \vec{s} is linked to \vec{t} , it follows that $|\vec{x} \vee \vec{s}| \leq |\vec{x}|$. Then, since $\vec{x} \vee \vec{s} \geq \vec{s}$ and \vec{r} is linked to \vec{s} , it follows that $|(\vec{x} \vee \vec{s}) \vee \vec{r}| \leq |\vec{x} \vee \vec{s}| \leq |\vec{x}|$. However, since $\vec{s} \leq \vec{r}$, $(\vec{x} \vee \vec{s}) \vee \vec{r} = \vec{x} \vee \vec{r}$, and so \vec{r} is linked to \vec{t} .

To see that the fourth condition is satisfied let us consider \overleftarrow{y}_j for some $j \in [m]$. There is some sequence of separations $\overleftarrow{x}_j = \overleftarrow{v}_0 \geq \overleftarrow{v}_1 \geq \dots \geq \overleftarrow{v}_t = \overleftarrow{u}_j$ that are the values \overleftarrow{y}_j takes during this process, corresponding to the t times we applied Lemma 15 to a pair containing the separation \overleftarrow{y}_j . Suppose that the other separations in those pairs were $\overleftarrow{y}_{i_1}, \overleftarrow{y}_{i_2}, \dots, \overleftarrow{y}_{i_t}$, and let us denote by \overleftarrow{w}_k the value of the separations \overleftarrow{y}_{i_k} at the time which we applied Lemma 15 to the pair $\{\overleftarrow{y}_j, \overleftarrow{y}_{i_k}\}$.

We claim inductively that for all $0 \leq r \leq t$

$$\overleftarrow{v}_0 \bigwedge_{k=1}^r \overleftarrow{w}_k \leq \overleftarrow{v}_r \leq \overleftarrow{v}_0.$$

The statement clearly holds for $r = 0$. Suppose it holds for $r - 1$. We obtain \overleftarrow{v}_r by applying Lemma 15 to the pair $\{\overleftarrow{v}_{r-1}, \overleftarrow{w}_r\}$, giving us the pair $\{\overleftarrow{v}_r, \overleftarrow{z}\}$. We have that $\overleftarrow{v}_{r-1} \wedge \overleftarrow{z} = \overleftarrow{v}_r$ and so, since $\overleftarrow{w}_r \leq \overleftarrow{z}$ it follows that

$$\overleftarrow{v}_{r-1} \wedge \overleftarrow{w}_r \leq \overleftarrow{v}_r \leq \overleftarrow{v}_{r-1}.$$

By the induction hypothesis we know that

$$\overleftarrow{v}_0 \bigwedge_{k=1}^{r-1} \overleftarrow{w}_k \leq \overleftarrow{v}_{r-1} \leq \overleftarrow{v}_0,$$

and so

$$\overleftarrow{v}_0 \bigwedge_{k=1}^r \overrightarrow{w}_k \leq \overleftarrow{v}_{r-1} \wedge \overrightarrow{w}_r \leq \overleftarrow{v}_r \leq \overleftarrow{v}_{r-1} \leq \overleftarrow{v}_0$$

as claimed.

For each of the \overleftarrow{w}_k there is some separation \overleftarrow{s}_k from our original set (that is some \overleftarrow{r}_i or \overleftarrow{x}_j) such that $\overleftarrow{w}_k \leq \overleftarrow{s}_k$ and so, since $\overrightarrow{s}_k \leq \overrightarrow{w}_k$, and since we apply Lemma 15 to each pair of separations in our original set, we have that

$$\overleftarrow{v}_0 \bigwedge_i \overrightarrow{r}_i \bigwedge_{k \neq j} \overrightarrow{x}_j \leq \overleftarrow{v}_0 \bigwedge_{k=1}^t \overrightarrow{w}_k.$$

So, recalling that $\overleftarrow{v}_0 = \overleftarrow{x}_j$ and $\overleftarrow{v}_t = \overleftarrow{u}_j$, we see that

$$\overleftarrow{x}_j \bigwedge_i \overrightarrow{r}_i \bigwedge_{k \neq j} \overrightarrow{x}_k \leq \overleftarrow{u}_j \leq \overleftarrow{x}_j$$

as claimed.

Finally we note that, if we apply Lemma 15 to a pair of separations (C, D) and (E, F) , resulting in the nested pair $\{(C', D'), (E', F')\}$, then

$$C \cup E = C' \cup E'.$$

Indeed, we have that $(C \cap F', D \cup E') = (C', D')$ and $(E \cap D', F \cup C') = (E', F')$ and so we have that $C' \cup E' = (C \cap F') \cup E' \supseteq C$ and similarly $C' \cup E' = C' \cup (E \cap D') \supseteq E$ and so $C' \cup E' \supseteq C \cup E$. However, since $C' \subseteq C$ and $E' \subseteq E$ we also have $C' \cup E' \subseteq C \cup E$. \square

4.2 Refining the essential parts

The content of Lemma 17 can be thought of as a procedure for turning an arbitrary set of separations into a star which is in some way ‘close’ to the original set, and is linked pairwise to the original set. We note that the second property guarantees us that this star lies in the same S_k as the original set.

Let us say a few words about the other properties of the star which represent this closeness. It will be useful to think about these properties in terms of how we can use this lemma to refine further an essential part in a k -tangle-distinguishing tree-decomposition.

Suppose $\{\overleftarrow{x}_j : j \in [m]\} = \mathcal{M}_I(O)$ for some k -tangle O , and $\{\overleftarrow{r}_i : i \in [n]\}$ is the star of separations at the vertex where O is contained in a tree-decomposition, specifically one where each r_i distinguishes efficiently some pair of k -tangles. By applying Lemma 17 we get a star $\{\overleftarrow{u}_j : j \in [m]\}$ satisfying the conclusions of the lemma. For each non-trivial \overleftarrow{x}_j , by Lemma 13, there exists an irredundant S_k -tree over $\mathcal{T}_k^* \cup \{\overleftarrow{x}_j\}$ containing \overleftarrow{x}_j as a leaf separation, such that \overleftarrow{x}_j is not the image of any other edge. We can then use Lemma 4 to shift each of these S_k -trees onto \overrightarrow{u}_j , giving us an S_k -tree over $\mathcal{T}_k^* \cup \{\overleftarrow{u}_j\}$. If \overleftarrow{x}_j is trivial then so is \overleftarrow{u}_j , and so there is an obvious S_k -tree over $\mathcal{T}_k^* \cup \{\overleftarrow{u}_j\}$ containing \overrightarrow{u}_j as a leaf separation, that with a single edge corresponding to u_j .

Doing the same for each k -tangle in the graph and taking the union all of these S_k -trees, together with the tree-decomposition from Corollary 9, will give us a refinement of this tree-decomposition which maintains the property of each inessential part being too small to contain a k -tangle, but also further refines the essential parts. The properties of the star given by Lemma 17 give us some measurement of how effective this process is in refining the essential parts of the graph.

We first note that, given a k -tangle O , which is contained in some part V_t of a k -tangle-distinguishing tree-decomposition, by the fifth property in Lemma 17 every vertex in the part V_t which lies on the small side of some maximal inessential separation in O will be in some inessential part of this refinement.

However this property is also satisfied by the rather naive refinement formed by just taking the union of some small separations (A_i, V) with the A_i covering the same vertex set. The problem with this naive decomposition is it does not really refine the part V_t , since there is still a part with vertex set V_t in the new decomposition. Ideally we would like our refinement to make this essential part as small as possible, to more precisely exhibit how the k -tangle O lies in the graph.

Our refinement comes some way towards this, as evidenced by the fourth condition. For example if we have some separation $\overleftarrow{s} = (A, B)$ which lies ‘behind’ some maximal inessential separation in O ,

that is $\overleftarrow{s} \leq \overleftarrow{x}_j$ for some j , and is nested $\mathcal{M}_I(O) \cup \{\overleftarrow{r}_i : i \in [n]\}$, then it is easy to check that the fourth property guarantees it will also lie behind some \overleftarrow{u}_k given by Lemma 17. So, in the refined tree-decomposition, the part containing O will not contain any vertices that lie strictly in the small side of such a separation, $A \setminus B$.

Suppose $\{\overleftarrow{r}_i : i \in [n]\}$ is an essential part in a tree-decomposition (T, \mathcal{V}) containing a tangle O . We say a vertex $v \in V$ is *inessentially separated from O relative to (T, \mathcal{V})* if there is a separation (A, B) which is nested with $\mathcal{M}_I(O) \cup \{\overleftarrow{r}_i : i \in [n]\}$ such that $v \in A \setminus B$, and there exists some $(X, Y) \in \mathcal{M}_I(O)$ such that $(A, B) \leq (X, Y)$. For example, in Figure 2, the vertices in the long paths are inessentially separated from the tangles corresponding to the complete subgraphs relative to the canonical tangle-distinguishing tree-decomposition.

Theorem 18. *For every graph G and $k \geq 3$ there exists a tree-decomposition (T, \mathcal{V}) of G of adhesion $< k$ with the following properties*

- *The tree-decomposition (T', \mathcal{V}') induced by the essential separations is canonical and distinguishes every k -tangle in G ;*
- *The torso of every inessential part has branch-width $< k$.*
- *For every essential part V_t which contains a tangle O , there are no vertices $v \in V_t$ which are inessentially separated from O relative to (T', \mathcal{V}') .*

Given a vertex $v \in V$ we say that x is *well separated from O* if there is a separation (A, B) which is nested with $\mathcal{M}(O)$ such that $v \in A \setminus B$, and there exists some $(X, Y) \in \mathcal{M}(O)$ such that $(A, B) \leq (X, Y)$.

We can think of the vertices which are well separated from O as being ‘far away’ from O in the graph. Indeed, if $\mathcal{M}(O)$ is a star, then O is a k -block, and the set of vertices well separated from O are just the vertices not in the k -block. In general a tangle will not correspond as closely to a concrete set of vertices as a k -block, and crossing separations in $\mathcal{M}(O)$ somehow demonstrate the uncertainty of whether a vertex ‘lives in’ O or not. However, if a separation $(A, B) \in O$ is nested with $\mathcal{M}(O)$, then O should be in some way fully contained in B , and so the vertices in $A \setminus B$ are ‘far away’ from O .

Question 19. *For every graph G , does there exist a tree-decomposition which distinguishes the k -tangles in a graph, whose essential parts are small in the sense that for each k -tangle O , there is no vertex x which can be well separated from O in the part of the tree-decomposition which contains O ? Does there exist such a tree-decomposition with the further property that the inessential parts have branch-width $< k$?*

References

- [1] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs i. existence and algorithms. *Journal of Combinatorial Theory, Series B*, 116:1–24, 2016.
- [2] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs ii. essential parts. *Journal of Combinatorial Theory, Series B*, 118:268–283, 2016.
- [3] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. Connectivity and tree structure in finite graphs. *Combinatorica*, 34(1):1–35, 2014.
- [4] R. Diestel, P. Eberenz, and J. Erde. Duality theorems for tangles in graphs. *arXiv preprint arXiv:1605.09139*, 2016.
- [5] R. Diestel, F. Hundertmark, and S. Lemanczyk. Profiles of separations: in graphs, matroids and beyond. *arXiv preprint arXiv:1110.6207*, 2016.
- [6] R. Diestel and S. Oum. Unifying duality theorems for width parameters, II. General duality. Preprint 2014.
- [7] R. Diestel and S. Oum. Tangle-tree duality in abstract separation systems. *arXiv preprint arXiv:1701.02509*, 2014.

- [8] R. Diestel and S. Oum. Tangle-tree duality: in graphs, matroids and beyond. *arXiv preprint arXiv:1701.02651*, 2014.
- [9] N. Robertson and P. D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B*, 52:153–190, 1991.
- [10] P. Seymour and R. Thomas. Graph searching and a min-max theorem for tree-width. *J. Combin. Theory (Series B)*, 58(1):22 – 33, 1993.
- [11] W. T. Tutte. *Graph Theory*. Addison-Wesley, 1984.