Innov. Graph Theory **3**, 2026, pp. 171–206
https://doi.org/10.5802/igt.21

INNOVATIONS IN GRAPH THEORY

OPTIMAL TREES OF TANGLES: REFINING THE ESSENTIAL PARTS

by Sandra ALBRECHTSEN

Abstract. — We combine the two fundamental fixed-order tangle theorems of Robertson and Seymour into a single theorem that implies both, in a best possible way. We show that, for every $k \in \mathbb{N}$, every tree-decomposition of a graph G which efficiently distinguishes all its k-tangles can be refined to a tree-decomposition whose parts are either too small to be home to a k-tangle, or as small as possible while being home to a k-tangle.

1. Introduction

Tangles were introduced by Robertson and Seymour [17] as a way to capture "highly cohesive" structures of a graph. Formally, a *k-tangle* in a graph G, for some $k \in \mathbb{N}$, is a consistent orientation of all the separations $\{A, B\}$ of G of order less than k, as (A, B) say, such that no three such oriented separations together cover the whole graph by the subgraphs induced on their "small sides" A. The idea is that every highly cohesive substructure of G will lie mostly on one side of every low-order separation, and therefore orients it towards that side, thereby defining a tangle.

Since its first introduction, the notion of a tangle and its framework in graphs have been generalized to so-called *abstract separation systems* [8]. Although these abstract separation systems are significantly more general than the separation system of a graph, the two fundamental theorems about graph tangles (the "tree-of-tangles" theorem and the "tangle-tree duality" theorem of Robertson and Seymour [17]) are still valid in this setting. We shall also prove most of our results in this more general setting.

Keywords: Tree of tangles, tree-decomposition, tangle-tree duality, abstract separation system.
2020 *Mathematics Subject Classification:* 05C83, 05C40, 06A07.

In what follows we assume familiarity only with graph tangles as described in [7, Chapter 12.5].

The *tree-of-tangles theorem* asserts the existence of a tree-decomposition $(T, \mathcal{V})$ of G, with $\mathcal{V} = (V_t)_{t \in T}$, say, which *distinguishes* all the k-tangles in G: for every pair of k-tangles there exists an edge of T which induces a separation that is oriented differently by these two k-tangles. For this, recall that G reflects the separation properties of T: similar as the deletion of an edge $e = \{t_1, t_2\} \in E(T)$ separates T into two subtrees $T_1 \ni t_1$ and $T_2 \ni t_2$, the sets $U_i := \bigcup_{t \in T_i} V_t$, for $i = 1, 2$, form a separation $\{U_1, U_2\}$ of G.

Since $(T, \mathcal{V})$ distinguishes all the k-tangles in G, each k-tangle *lives* in a different part: the part $G[V_t]$ specified by the property that the given k-tangle orients all the separations induced by edges of T incident with t to that side which contains V_t. We call the set of these oriented separations the *star associated with* V_t. Such a tangle-distinguishing tree-decomposition provides information about the overall structure of the graph and the location of the tangles inside it. A part $V_t \in \mathcal{V}$ is called *essential* if there is a tangle living in it, and otherwise *inessential*.

In general, the inessential parts might contain a large portion of G. The tree-decomposition then tells us nothing about the structure of this portion. However, if there are no k-tangles in G at all, there is another theorem which does tell us something about its structure: the *tangle-tree duality theorem*. This asserts the existence of a tree-decomposition in which each part is too small to be home to a tangle, and which thus witnesses that G has no k-tangles (since any k-tangle would have to live in some part). Formally, such a tree-decomposition has the property that the stars associated with its parts are elements of $\mathcal{T}_k$: the collection of all those sets of at most three oriented separations which every k-tangle has to avoid.

As inessential parts are not home to any tangles, the tangle-tree duality theorem applies locally to these parts. We thus obtain tree-decompositions of all the inessential parts in $\mathcal{V}$ into smaller parts, too small to accommodate a k-tangle. Can all these tree-decompositions be combined into a single tree-decomposition of G that refines $(T, \mathcal{V})$?

In general, this will not be possible. However, Erde [16] showed that it can be done if those local tree-decompositions are constructed carefully for this purpose. The following theorem follows directly from [16, Lemma 3.1]:

Theorem 1.1 (Erde 2017 [16]). — *Let $(\widetilde{T}, \widetilde{\mathcal{V}})$ be a tree-decomposition of a graph G which distinguishes all the k-tangles in G, for some $k \geqslant 3$, and is such that every separation of G induced by an edge of $\widetilde{T}$ efficiently*

distinguishes some pair of k-tangles in G. Then there exists a tree-decomposition $(T, \mathcal{V})$ of G which refines $(\widetilde{T}, \widetilde{\mathcal{V}})$ and is such that every star associated with an inessential part of $(T, \mathcal{V})$ is a star in $\mathcal{T}_k$.

See Section 2 for definitions.

The procedure of refining $(T, \mathcal{V})$ by further decomposing its inessential parts can be seen as a way to describe the structure of those parts in a more precise way. However, just as the inessential parts of $(T, \mathcal{V})$ may contain a large portion of G, another large portion may be hidden inside its essential parts. The tree-decomposition $(T, \mathcal{V})$ does not tell us much about this portion. We know from the fact that these parts are essential that there is a tangle living there, but an essential part of $(T, \mathcal{V})$ can also contain many vertices and edges that have nothing to do with that tangle.

To get a better insight into the structure of the essential parts of $(T, \mathcal{V})$ too, it would be desirable to decompose them by further tree-decompositions as well. Overall, it would be useful to have a single tree-decomposition of G whose inessential and essential parts are both as small as possible. As every essential part is home to a tangle, by definition, and every tangle corresponds to some highly cohesive (but maybe fuzzy) substructure of G which cannot be divided into small parts, "small" can only be achieved relatively to the tangle living in that part. In other words, the essential parts should be as small as possible so that they can still be home to a tangle.

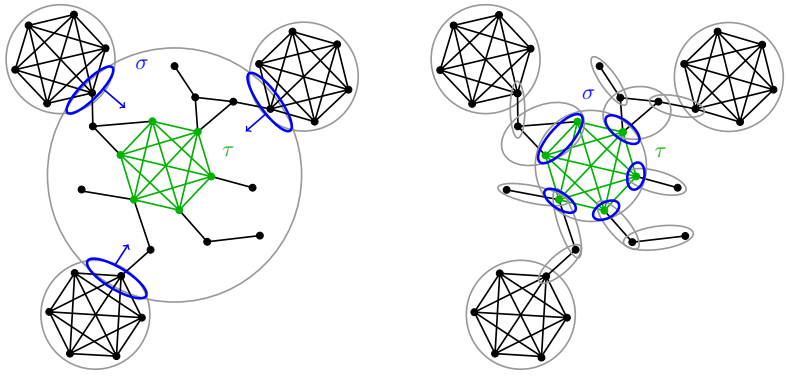

Figure 1.1. A tree-decomposition of G that distinguishes all its tangles (left), and a refinement of its central essential part (right).

For example, consider the graph G in Figure 1.1. It contains a central K_6, which induces a 3-tangle τ of G: the set of all separations of G of order at most 2 oriented towards the central K_6. The star of separations $\sigma \subseteq \tau$

indicated in the figure on the left side, which distinguishes τ from the other 3-tangles of G (those induced by the other three K_6 subgraphs), lies far away from the central K_6, the "essence" of τ. This results in an essential part which contains a lot more vertices than those contained in the central K_6. But although τ lives in that part, those other vertices do not really "belong to τ".

In the tree-decomposition on the right side of Figure 1.1 the star σ has been moved closer to the core of τ, the K_6 that induces it. In this decomposition of G, every essential part contains only the vertices of its K_6, while every inessential part is small in that it contains only few vertices. Clearly, the tree-decomposition in on the right side of Figure 1.1 captures the structure of G in a more precise way than the tree-decomposition in Section 1 does, and thus provides more information about the graph.

Erde [16] already introduced a method to refine the essential parts of a given tree-decomposition so that the newly arising essential parts do not contain vertices that are "inessentially separated" from the tangle living in that part. This result comes somewhat close to our goal of decreasing the size of the essential parts by refining them. However, it is not strong enough to decrease the size of the essential parts as much as possible, so that even in the example described above, where the tangle τ is induced by a clique, there are several vertices not in the clique which are not "inessentially separated" from τ. Thus, if we apply the method from [16] to the essential part of the tree-decomposition on the left side of Figure 1.1 that accommodates τ, then the central essential part of the refined tree-decomposition will still contain more vertices than just those in the clique.

In this paper we show that the essential parts of tangle-distinguishing tree-decompositions can be refined in the best possible way: so that the newly arising essential parts contain as few vertices as possible while being home to a tangle (see Section 2 for definitions):

THEOREM 1.2. — *Let $(\widetilde{T}, \widetilde{\mathcal{V}})$ be a tree-decomposition of a graph G which distinguishes all the k-tangles in G, for some $k \in \mathbb{N}$ so that every separation induced by an edge of $\widetilde{T}$ efficiently distinguishes some pair of k-tangles in G. Then there exists a tree-decomposition $(T, \mathcal{V})$ of G which refines $(\widetilde{T}, \widetilde{\mathcal{V}})$ and is such that*

(i) every star associated with an inessential part of $(T, \mathcal{V})$ is a star in $\mathcal{T}_k$;

(ii) every essential part V_t of $(T, \mathcal{V})$ is of smallest size among all the exclusive parts of tree-decompositions of G that are home to the k-tangle living in V_t.

Additionally, the tree-decomposition $(T, \mathcal{V})$ from Theorem 1.2 has the following three properties. First, each of its inessential parts has size $\leqslant 3k-3$. Second, the unique k-tangle τ that lives in an essential part V_t of size $> 3k-3$ is induced by V_t, in the sense that a separation $\{A, B\}$ of G of order less than k is oriented by τ as (A, B) if and only if $|A \cap V_t| < |B \cap V_t|$. Third, the restriction of τ to $G[V_t]$ is a k-tangle of $G[V_t]$; thus, $G[V_t]$ alone already witnesses that τ is indeed a k-tangle in G.

If a tangle τ in G is induced by a separable k-block in G (see Section 4), then the essential part of our tree-decomposition that is home to τ will be equal to that k-block. Thus, the tree-decomposition $(T, \mathcal{V})$ of G isolates all the separable k-blocks that induce k-tangles, in that they appear as parts of the tree-decomposition. See Theorem 4.10 for details.

We also prove a variant of Theorem 1.2 for abstract separation systems that are not induced by a graph. For this, recall that a tree-decomposition of a graph induces a "nested" set of separations of that graph and vice versa. Moreover, one can generalize the notion of a k-tangle to that of an $\mathcal{F}$-tangle by allowing other sets $\mathcal{F}$, instead of $\mathcal{T}_k$, whose elements an $\mathcal{F}$-tangle needs to avoid. We will in fact prove the following more general version of Theorem 1.2 (see Section 2 for definitions):

THEOREM 1.2′. — *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a friendly set of stars in $\vec{S}_k(G)$. Further, let $\widetilde{N} \subseteq S_k(G)$ be a nested set of separations that distinguishes the all the $\mathcal{F}$-tangles of $S_k(G)$ so that every separation in $\widetilde{N}$ efficiently distinguishes a pair of $\mathcal{F}$-tangles of $S_k(G)$. Then there exists a nested set $N \subseteq S_k(G)$ with $\widetilde{N} \subseteq N$ such that*

- *(i) every inessential node of N is a star in $\mathcal{F}$;*
- *(ii) the interior of every essential node of N is of smallest size among all the exclusive stars contained in the $\mathcal{F}$-tangle living at that node.*

Theorem 1.2′ extends verbatim to abstract separation systems except for its condition (ii), whose reference to "size" involves vertices. Vertices need not exist in the setting of abstract separation systems, not even in the matroid setting. The following variant of Theorem 1.2′ adapts (ii) to this more general setting while leaving (i) unchanged. Note that the maximality in (ii) now refers to the nodes of N itself, which are sets of oriented separations contained in the relevant tangle. This maximality translates to the minimality of the interior of N referred to in (ii) of Theorem 1.2′, which in the abstract setting is no longer defined.

THEOREM 1.3. — *Let $\vec{S}$ be a submodular separation system inside some distributive universe and let*

$$\mathcal{F} \supseteq \big\{\{\vec{r}, \overleftarrow{s}\} \subseteq \vec{S} : \vec{r} \leqslant \vec{s} \text{ and } \vec{r} \vee \overleftarrow{s} \text{ is co-small}\big\}$$

be a friendly set of stars in $\vec{S}$. Further, let $\widetilde{N}$ be a nested set of good separations that distinguishes all the $\mathcal{F}$-tangles of S. Then there exists a nested set $N \subseteq S$ with $\widetilde{N} \subseteq N$ such that

(i) every inessential node of N is a star in $\mathcal{F}$;
(ii) every essential node of N is maximal in the $\mathcal{F}$-tangle living at it.

Theorem 1.3 requires that every separation in $\widetilde{N}$ is good rather than that it efficiently distinguishes two tangles as in Theorem 1.2. This is a necessary adjustment as the notion of "efficiency" does not extend to abstract separation systems. In separation systems induced by a graph these two properties agree.

Condition (ii) in Theorem 1.3 implies that one cannot refine the nested set N from Theorem 1.3 any further, so N is optimal as claimed. We remark that by [2, Theorem 5.8] there always exists a nested set $\widetilde{N}$ which satisfies the premise of Theorem 1.3.

This paper is organised as follows. In Section 2 we give a short introduction to abstract separation systems. We then recall the definitions and lemmas from [2] about refining inessential stars in Section 3. In Section 4 and 5 we prove Theorem 1.2 and 1.3, respectively.

2. Abstract separation systems and profiles

In this section we give a short overview of the important objects, definitions and theorems that we are going to use later. For a more detailed introduction to abstract separation systems and their tangles we refer the reader to [12] and [8]. The definitions are taken from [8, 9, 12, 13, 15]. We use the basic graph-theoretic notions from [7]. In particular, we may speak of a vertex $v \in G$ (rather than $v \in V(G)$).

2.1. Separation systems

A *separation* of a graph G is a set $\{A, B\}$ of subsets of $V(G)$ such that $A \cup B = V(G)$ and there is no edge in G between $A \setminus B$ and $B \setminus A$. The ordered pairs (A, B) and (B, A) of a separation $\{A, B\}$ are its *orientations*.

The set $\vec{U}(G)$ of all oriented separations of G is partially ordered by $(A,B) \leqslant (C,D)$ if $A \subseteq C$ and $D \subseteq B$. Inverting the orientations of a separation $\{A,B\}$ is an involution $(A,B) \mapsto (A,B)^* := (B,A)$ on $\vec{U}(G)$ which reverses the partial order, that is $(A,B) \leqslant (C,D)$ if and only if $(B,A) \geqslant (D,C)$.

Any two separations $(A,B),(C,D)$ of G have an *infimum* $(A \cap C, B \cup D)$ and a *supremum* $(A \cup C, B \cap D)$ in $\vec{U}(G)$. Hence, $\vec{U}(G)$ forms a lattice together with the operations $(A,B) \wedge (C,D) := (A \cap C, B \cup D)$ as well as $(A,B) \vee (C,D) := (A \cup C, B \cap D)$.

The number $|A \cap B|$ is the *order* of the separation $\{A,B\}$. It is easy two check that the orders of the infimum and supremum of any two separations $(A,B),(C,D)$ of G sum up to

$$|(A,B) \vee (C,D)| + |(A,B) \wedge (C,D)| = |(A,B)| + |(C,D)|.$$

The important part of this equality is that the left side is never larger than the right side, which is called *submodularity*.

A *separation system* $\vec{S} = (\vec{S}, \leqslant, ^*)$ is a partially ordered set $\vec{S}$ with an order-reversing involution *, i.e. if $\vec{s} \leqslant \vec{r} \in \vec{S}$, then $(\vec{s})^* \geqslant (\vec{r})^*$. We denote $(\vec{s})^*$ by $\overleftarrow{s}$ and $(\overleftarrow{s})^*$ by $\vec{s}$. The elements $\vec{s}, \overleftarrow{s} \in \vec{S}$ are called *oriented separations*. An *unoriented* separation is a set $s := \{\vec{s}, \overleftarrow{s}\}$ for some $\vec{s} \in \vec{S}$, and $\vec{s}$, $\overleftarrow{s}$ are the *orientations* of s. Note that there are no default orientations: Once we denoted one orientation by $\vec{s}$ the other one will be $\overleftarrow{s}$, and vice versa. The set of all sets $\{\vec{s}, \overleftarrow{s}\} \subseteq \vec{S}$ is denoted by S. We will use terms defined for unoriented separations also for oriented ones and vice versa if that is possible without causing ambiguities. Moreover, if the context is clear, we will simply refer to both oriented and unoriented separations as "separations".

Two unoriented separations $r, s \in S$ are *nested* if they have orientations that can be compared using $\leqslant$, otherwise they *cross*. If $r, s \in S$ are nested, then both r and $\vec{r}$ are *nested* with both of $s, \vec{s}$. Analogously, if $r, s \in S$ cross, then both $r, \vec{r}$ *cross* both of $s, \vec{s}$. A set of separations is *nested* if all its elements are pairwise nested. Two sets R, R' of separations are *nested* if every element of R is nested with every element of R'.

If a separation $s \in S$ satisfies $\vec{s} = \overleftarrow{s}$ we call s *degenerate*. A separation $\vec{s} \in \vec{S}$ is called *small* and $\overleftarrow{s}$ *co-small* if $\vec{s} \leqslant \overleftarrow{s}$. In the case of separations of a graph G these are exactly the separations of the form $(A, V(G))$. A separation $\vec{s} \in \vec{S}$ is *trivial* in $\vec{S}$ if there exists a separation $\vec{r} \in \vec{S}$ such that $\vec{s} < \vec{r}$ and $\vec{s} < \overleftarrow{r}$. Note that trivial separations are small. Conversely, all but the largest small separations (with respect to $\leqslant$) are trivial. A set $\mathcal{F} \subseteq 2^{\vec{S}}$

of subsets of $\vec{S}$ is *standard (for $\vec{S}$)* if $\{\overleftarrow{r}\} \in \mathcal{F}$ for every trivial separation $\vec{r} \in \vec{S}$.

A nested set $\sigma \subseteq \vec{S}$ without degenerate separations is called a *star* if for any $\vec{r}, \vec{s} \in \sigma$ it holds that $\vec{r} \leqslant \overleftarrow{s}$. If G is a graph and $\sigma \subseteq \vec{U}(G)$ a star, then $\mathrm{int}(\sigma) := \bigcap_{(A,B)\in\sigma} B$ is the *interior of σ*.

A star σ is called *proper* if for every distinct $\vec{r}, \vec{s} \in \sigma$ the relation $\vec{r} \leqslant \overleftarrow{s}$ is the only one, i.e. $\vec{r} \not\leqslant \vec{s}$, $\vec{r} \not\geqslant \vec{s}$ and $\vec{r} \not\geqslant \overleftarrow{s}$. A star σ is *proper in $\vec{S}$* if σ is proper and $\sigma \neq \{\overleftarrow{s}\}$ for a trivial separation $\vec{s} \in \vec{S}$.

The partial order on $\vec{S}$ also relates the proper stars in $\vec{S}$: if $\sigma, \tau \subseteq \vec{S}$ are proper stars, then $\sigma \leqslant \tau$ if and only if for every $\vec{s} \in \sigma$ there exists some $\vec{r} \in \tau$ such that $\vec{s} \leqslant \vec{r}$. Note that this relation is again a partial order [9]. A proper star in $\vec{S}$ is *maximal* if it is a maximal element in the set of all proper stars in $\vec{S}$ with the partial order defined above.

We call a separation system $(\vec{U}, \leqslant, ^*)$ a *universe* of separations, and denote it with $\vec{U} = (\vec{U}, \leqslant, ^*, \wedge, \vee)$, if it is a lattice, i.e. if any two separations $\vec{r}, \vec{s} \in \vec{U}$ have a *supremum* $\vec{r} \vee \vec{s}$ and an *infimum* $\vec{r} \wedge \vec{s}$ in $\vec{U}$.

A universe $\vec{U}$ of separations is *distributive* if it is distributive as a lattice, i.e. if

$$\vec{r} \vee (\vec{s} \wedge \vec{t}) = (\vec{r} \vee \vec{s}) \wedge (\vec{r} \vee \vec{t}) \text{ and } \vec{r} \wedge (\vec{s} \vee \vec{t}) = (\vec{r} \wedge \vec{s}) \vee (\vec{r} \wedge \vec{t})$$

for all $\vec{r}, \vec{s}, \vec{t} \in \vec{U}$.

Given two separations $r, s \in U$ we call the separations $\{(\vec{r} \vee \vec{s}), (\vec{r} \vee \vec{s})^*\}$, $\{(\vec{r} \wedge \vec{s}), (\vec{r} \wedge \vec{s})^*\}$, $\{(\overleftarrow{r} \vee \vec{s}), (\overleftarrow{r} \vee \vec{s})^*\}$ and $\{(\overleftarrow{r} \wedge \vec{s}), (\overleftarrow{r} \wedge \vec{s})^*\}$ their *corner separations*. A simple but quite useful observation about corner separations is the following:

Lemma 2.1 ([8, Lemma 3.2]). — *Let $r, s \in S$ be two crossing separations. Every separation t that is nested with both r and s is also nested with all four corner separations of r and s.*

A function $|\cdot| : \vec{U} \to \mathbb{R}_{\geqslant 0}$ is called an *order function* of $\vec{U}$ if $|\vec{s}| = |\overleftarrow{s}|$ for all $s \in U$. If $\vec{U}$ comes with an order function that is *submodular*, i.e.

$$|\vec{r} \vee \vec{s}| + |\vec{r} \wedge \vec{s}| \leqslant |\vec{r}| + |\vec{s}| \text{ for all } \vec{r}, \vec{s} \in \vec{U}$$

we call $\vec{U}$ a *submodular* universe. Then $|\vec{s}|$ is the *order* of $\vec{s}$ and s. For an integer $k > 0$, the induced sets $\vec{S}_k := \{\vec{s} \in \vec{U} : |\vec{s}| < k\}$ are separation systems on their own. Though they need not be universes, as in general $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ for two separations $\vec{r}, \vec{s} \in \vec{S}_k$ need not both lie in $\vec{S}_k$. Note that we take the infimum and supremum here with respect to $\vec{U}$. However, by the submodularity of $|\cdot|$, at least one of them has to be contained in $\vec{S}_k$.

This property is rather useful and motivates the following generalization of submodularity which can be applied to universes without making use of the external concept of an order function:

A separation system $\vec{S}$ is *inside* a universe $\vec{U}$ if $\vec{S} \subseteq \vec{U}$ and the partial order and the involution on $\vec{S}$ are the ones induced by $\vec{U}$. Further, a separation system $\vec{S}$ inside some universe $\vec{U}$ is called *submodular* if

$$\text{for all } \vec{r}, \vec{s} \in \vec{S} : \vec{r} \vee \vec{s} \in \vec{S} \text{ or } \vec{r} \wedge \vec{s} \in \vec{S}.$$

2.2. Profiles and tangles

An *orientation* of a set S of unoriented separations is a set $O \subseteq \vec{S}$ which contains every degenerate separation from $\vec{S}$, and exactly one orientation $\vec{s}$ or $\overleftarrow{s}$ of every non-degenerate separation in S. A subset $O \subseteq \vec{S}$ is *consistent* if it does not contain both $\overleftarrow{r}$ and $\vec{s}$ whenever $\vec{r} < \vec{s}$ for distinct $r, s \in S$. If $\mathcal{O}$ is a set of consistent orientations of $\vec{S}$, then we call a star $\sigma \subseteq \vec{S}$ *essential (for $\mathcal{O}$)* if $\sigma \subseteq O$ for some $O \in \mathcal{O}$. Otherwise σ is called *inessential (for $\mathcal{O}$)*.

A star $\sigma \subseteq \vec{S}$ is *exclusive* for some set $\mathcal{O}$ of consistent orientations of S if it is contained in exactly one orientation in $\mathcal{O}$. If $O \in \mathcal{O}$ is that orientation, we say that σ is *O-exclusive (for $\mathcal{O}$)*.

A non-degenerate separation $s \in S$ *distinguishes* two orientations O_1 and O_2 of S if O_1 and O_2 orient s differently. A set of separations $N \subseteq S$ *distinguishes* a set $\mathcal{O}$ of orientations if any two distinct orientations in $\mathcal{O}$ are distinguished by some separation in N. If S lies inside a submodular universe, a separation distinguishes a pair of orientations O_1 and O_2 of S *efficiently* if it distinguishes them and O_1 and O_2 cannot be distinguished by any separation of lower order.

Let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a set of subsets of $\vec{S}$. An orientation P of S is an *$\mathcal{F}$-tangle (of S)* if P is consistent and does not contain any element of $\mathcal{F}$ as a subset.

If $\vec{S}$ is a separation system inside some universe $\vec{U}$, we call a consistent orientation P of S a *profile of S* if it satisfies that

$$\text{for all } \vec{r}, \vec{s} \in P \text{ the separation } (\vec{r} \vee \vec{s})^* \text{ does not lie in } P.$$

A profile of S which contains no co-small separation is *regular*. For the special case of separations of a graph G, a *k-profile* is a profile of $S_k(G)$.

We can describe the profiles of S as $\mathcal{F}$-tangles as follows. Set $\mathcal{P}_S := \{\{\vec{r}, \vec{s}, (\vec{r} \vee \vec{s})^*\} : \vec{r}, \vec{s} \in \vec{S}\} \cap 2^{\vec{S}}$. Then the set of all $\mathcal{P}_S$-tangles of S is exactly the set of all profiles of S. We say that $\mathcal{F}$ is *profile-respecting (for $\vec{S}$)* if every $\mathcal{F}$-tangle of S is a profile of S.

The profiles which we consider in this paper will typically be $\mathcal{F}$-tangles for a standard and profile-respecting set of stars which contains $\{\overleftarrow{r}\}$ for every small $\vec{r} \in \vec{S}$, and thus these profiles will be regular.

For the special case of separations of a graph G, a *k-tangle in G* is a $\mathcal{T}_k$-tangle of $S_k(G)$ where

$$\mathcal{T}_k := \{\{(A_i, B_i) : i \in [3]\} \subseteq \vec{S}_k(G) : G[A_1] \cup G[A_2] \cup G[A_3] = G\}.$$

2.3. Tree sets, S-trees and tree-decompositions

A *tree set* is a nested separation system without degenerate or trivial separations. A star $\sigma \subseteq \overrightarrow{N}$ is called a *node of* N if there is a consistent orientation O of N such that σ is the set of maximal elements in O. (In [9] the nodes of N are called "splitting stars".) If $\vec{S}$ is a separation system, and if $N \subseteq S$ consists only of separations which each distinguish some pair of consistent orientations of S, then $(\overrightarrow{N}, \leqslant, ^*)$ is a tree set where $\leqslant$ and * are the ones induced by $\vec{S}$. Since all relevant nested sets in this paper will be of this form, we will use the term "node" also for nested sets.

Let $\vec{S}$ be again a separation system. An *S-tree* is a pair (T, α) of a (graph-theoretical) tree T and a map $\alpha : \vec{E}(T) \to \vec{S}$ from the oriented edges $\vec{E}(T)$ of T to $\vec{S}$ such that $\alpha(\overleftarrow{e}) = \overleftarrow{s}$ if $\alpha(\vec{e}) = \vec{s}$. If $x \in V(T)$ is a leaf of T and $t \in V(T)$ its unique neighbour, then we call $\alpha(x, t) \in \vec{S}$ a *leaf separation (of T)*.

An S-tree (T, α) is *over* a set $\mathcal{F} \subseteq 2^{\vec{S}}$ if, for every node $t \in T$, we have $\{\alpha(t', t) : (t', t) \in \vec{E}(T)\} \in \mathcal{F}$. If $\mathcal{F}$ is a set of stars in $\vec{S}$, then we call $\sigma_t := \{\alpha(t', t) : (t', t) \in \vec{E}(T)\} \subseteq \vec{S}$ for a node $t \in T$ the *star associated with t (in T)*.

Every S-tree over a set of stars *induces* a tree set $\overrightarrow{N} := \mathrm{im}(\alpha)$ via α if no $\alpha(e)$ is trivial or degenerate for some edge $e \in \vec{E}(T)$. Conversely, if N is a *regular* tree set in S, i.e. $\overrightarrow{N}$ does not contain any small separations, then one can obtain an S-tree (T, α) as follows. We take the set of all nodes of $\overrightarrow{N}$ as the vertex set of T and N as the edges of T where we let a separation $s \in N$ be incident to the two nodes of $\overrightarrow{N}$ that contain $\vec{s}$ and $\overleftarrow{s}$, respectively.

THEOREM 2.2 ([9, Theorem 6.9]). — *Let $\vec{S}$ be a separation system and $N \subseteq S$ a regular tree set. Then there exists an S-tree (T, α) with $\mathrm{im}(\alpha) = \overrightarrow{N}$ such that the stars associated with nodes of T are precisely the nodes of $\overrightarrow{N}$.*

This motivates the name "nodes" for the splitting stars of N. It is shown in [9] that the S-tree from Theorem 2.2 is unique up to isomorphisms. Therefore, we say that N *induces* the S-tree (T, α).

We say that a consistent orientation O *lives* at a node σ of N (or equivalently σ *is home* to O) if $\sigma \subseteq O$. Similarly, we say that O lives at a node t of an S-tree (T, α) if $\sigma_t \subseteq O$. It is easy to see that every consistent orientation of S has to live at a (unique) node of any regular tree set N.

Given a set $\mathcal{O}$ of consistent orientations of S and an S-tree (T, α), we call a node $t \in T$ *essential (for $\mathcal{O}$)* if there is an orientation in $\mathcal{O}$ which lives at t and otherwise *inessential (for $\mathcal{O}$)*.

A *tree-decomposition* of a graph G is a pair $(T, \mathcal{V})$ of a tree T together with a family $\mathcal{V} = (V_t)_{t \in T}$ of subsets of $V(G)$ such that $\bigcup_{t \in T} G[V_t] = G$, and such that for every vertex $v \in G$, the graph $T[\{t \in T : v \in V_t\}]$ is connected. The sets $V_t \in \mathcal{V}$ are called the *parts* of $(T, \mathcal{V})$. The number $\max\{|V_t \cap V_{t'}| : \{t, t'\} \in E(T)\}$ is the *adhesion* of $(T, \mathcal{V})$.

Every (oriented) edge $(t_1, t_2) \in \vec{E}(T)$ *induces* a separation $\vec{s}_{(t_1,t_2)} := (\bigcup_{t \in T_1} V_t, \bigcup_{t \in T_2} V_t)$ of G where $T_1 \ni t_1$ and $T_2 \ni t_2$ are the two components of $T - (t_1, t_2)$ [7]. It is easy to check that the set of separations induced by the edges of a tree-decomposition is nested. Conversely, it follows from Theorem 2.2 that every regular nested set N of separations of a graph G is induced by the edges of some tree-decomposition $(T, \mathcal{V})$ of G. In fact, if we require the tree T to be minimal, then there is a unique such tree-decomposition (up to isomorphisms). We say that N *induces* this tree-decomposition. In particular, if $(T, \mathcal{V})$ is induced by N, then the *stars* $\sigma_t := \{\vec{s}_{(u,t)} : (u, t) \in \vec{E}(T)\}$ *associated with the parts* V_t of $(T, \mathcal{V})$ are precisely the nodes of N.

Similar as to nested sets, we say that a consistent orientation O of $S_k(G)$, for some $k \in \mathbb{N}$, *lives* in a part $V_t \in \mathcal{V}$ if $\sigma_t \subseteq O$. Further, given a set $\mathcal{O}$ of consistent orientations of $S_k(G)$, we call V_t *essential (for $\mathcal{O}$)* if there is an orientation in $\mathcal{O}$ that lives in V_t and otherwise *inessential (for $\mathcal{O}$)*.

Given two tree-decompositions $(\widetilde{T}, \widetilde{\mathcal{V}})$ and $(T, \mathcal{V})$ of G, we say that $(T, \mathcal{V})$ *refines* $(\widetilde{T}, \widetilde{\mathcal{V}})$ if the set N of all the separations induced by the edges of T *refines* the set $\widetilde{N}$ induced by the edges of $\widetilde{T}$ in that $\widetilde{N} \subseteq N$.

2.4. Friendly sets of stars

Let $\vec{S}$ be a separation system inside some universe, and let $\vec{r}$ be a separation in $\vec{S}$ which is neither degenerate nor trivial, and set $S_{\geqslant \vec{r}} := \{x \in S : \vec{x} \geqslant \vec{r} \text{ or } \overleftarrow{x} \geqslant \vec{r}\}$. We say that $\vec{s}$ *emulates* $\vec{r}$ *in* $\vec{S}$ if $\vec{s} \geqslant \vec{r}$ and for every

$\vec{x} \in \vec{S}$ with $\vec{x} \geqslant \vec{r}$ we have $\vec{s} \vee \vec{x} \in \vec{S}$. We can then define a function

$$f\downarrow_{\vec{s}}^{\vec{r}} : \vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\} \to \vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\} \text{ by}$$
$$f\downarrow_{\vec{s}}^{\vec{r}}(\vec{x}) := \vec{x} \vee \vec{s} \text{ and } f\downarrow_{\vec{s}}^{\vec{r}}(\overleftarrow{x}) := (\vec{x} \vee \vec{s})^* \text{ for } \vec{r} \leqslant \vec{x}.$$

If (T, α) is an S-tree and $\vec{s}$ emulates $\vec{r}$, we can define $\alpha' := f\downarrow_{\vec{s}}^{\vec{r}} \circ \alpha$. It is then easy to see that (T, α') is again an S-tree, and we call (T, α') the *shift of* (T, α) *onto* $\vec{s}$.

Given a set of stars $\mathcal{F} \subseteq 2^{\vec{S}}$ we say that $\vec{s}$ *emulates* $\vec{r}$ *in* $\vec{S}$ *for* $\mathcal{F}$ if $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$ and for every star $\sigma \in \mathcal{F}$ with $\sigma \subseteq \vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\}$ that contains an element $\vec{t} \geqslant \vec{r}$ it holds that $f\downarrow_{\vec{s}}^{\vec{r}}(\sigma) \in \mathcal{F}$.

A set $\mathcal{F}$ of stars in $\vec{S}$ is *closed under shifting* if whenever $\vec{s} \in \vec{S}$ emulates some $\vec{r} \in \vec{S}$, then it also emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$. Further, a set $\mathcal{F}$ of stars in $\vec{S}$ is *friendly* if $\mathcal{F}$ is standard, profile-respecting, closed under shifting and contains $\{\overleftarrow{r}\}$ for every small separation $\vec{r} \in \vec{S}$.

The condition that a set of stars is friendly might seem to be rather strong at first but in practice it can typically be satisfied. Diestel, Eberenz and Erde [10] showed that any set $\mathcal{F} \subseteq 2^{\vec{S}}$ can be transformed into a standard set $\widehat{\mathcal{F}}$ of stars which is closed under shifting so that an orientation of S is a regular $\widehat{\mathcal{F}}$-tangle if and only if it is a regular $\mathcal{F}$-tangle.

Lemma 2.3 ([10, Lemmas 11 & 14]). — *Let $\vec{S}$ be a submodular separation system inside some universe and let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a standard set. Then there exists a standard set of stars $\widehat{\mathcal{F}} \subseteq 2^{\vec{S}}$ that is closed under shifting such that an orientation $\mathcal{O}$ of S is a regular $\mathcal{F}$-tangle if and only if it is a regular $\widehat{\mathcal{F}}$-tangle.*

Especially, Lemma 2.3 implies that $\mathcal{P}_S$ can be transformed into a set $\widehat{\mathcal{P}}_S$ of stars which is closed under shifting such that the set of all $\widehat{\mathcal{P}}_S$-tangles of S is precisely the set of all profiles of S. Thus, any set $\mathcal{F} \subseteq 2^{\vec{S}}$ can be transformed into a friendly set $\overline{\mathcal{F}} := \widehat{\mathcal{F}} \cup \widehat{\mathcal{P}}_S \cup \{\{\overleftarrow{s}\} : \vec{s} \in \vec{S} \text{ is trivial}\}$ of stars in $\vec{S}$ so that an orientation of S is an $\overline{\mathcal{F}}$-tangle if and only if it is a regular profile and an $\mathcal{F}$-tangle of S.

Moreover, for the special case of a separation system $S_k(G)$ that comes from a graph G, it is easy to check that the set $\mathcal{T}_k^*$ that consists of all stars in $\mathcal{T}_k$ is a friendly set of stars. Further, Diestel and Oum have shown that the $\mathcal{T}_k^*$-tangles of $S_k(G)$ are precisely the k-tangles in G if $|G| \geqslant k$ [14, Lemma 4.2].

3. Refining inessential stars

The idea of refining tangle-distinguishing tree-decompositions has its origin in [16]. There, Erde proved that in separation systems of the form S_k inside submodular universes one can, under certain circumstances, refine the inessential nodes of tangle-distinguishing nested sets of separations so that all the inessential nodes of the refinement are "small".

Since then, this theorem has been generalized to abstract separation systems [2]. Much of the theory that was introduced in [2], including the main result which generalizes Theorem 1.1, will play an important role in the proofs of Theorem 1.2 and 1.3. Therefore, in this section we recap those definitions and lemmas.

Let P be a profile of a separation system S inside some universe, and let $\vec{s} \in \vec{S}$ be a separation. We say that $\vec{s}$ is *closely related to* P if $\vec{s} \in P$ and $\vec{r} \wedge \vec{s} \in \vec{S}$ for every $\vec{r} \in P$. Further, we call a subset of $\vec{S}$ *closely related* to P if all its elements are closely related to P.

A separation $\vec{r} \in P$ is *maximal* in P if it is a maximal element in P with respect to the partial order on P induced by $\vec{S}$. An example of separations which are closely related to P are those which are maximal in P. In fact, even more is true:

Proposition 3.1 ([2, Lemma 4.4]). — *Let $\vec{S}$ be a submodular separation system, and let P be a profile of S. Further, let $Y \subseteq P$ be a nested set such that the inverse $\overleftarrow{y}$ of every $\vec{y} \in Y$ is closely related to some profile of S, and let $P_Y \subseteq P$ be the set of all separations in P that are nested with Y. Then every maximal separation in P_Y is closely related to P.*

The next observation, which we will use throughout this paper, describes a sufficient condition for a separation to be closely related to some profile:

Proposition 3.2 ([2, Proposition 3.3]). — *Let $\vec{S}$ be a submodular separation system, and let $\vec{s} \in \vec{S}$ be closely related to some profile P of S. Further, let $\vec{r} \in \vec{S}$ with $\vec{r} \leqslant \vec{s}$, and suppose that $\vec{r} \wedge \vec{u} \in \vec{S}$ for every $\vec{u} \leqslant \vec{s}$. Then $\vec{r}$ is closely related to P.*

Moreover, infima of separations which are closely related to some profile are, under certain circumstances, again closely related to that profile:

Proposition 3.3 ([2, Proposition 4.3]). — *Let $\vec{S}$ be a submodular separation system and $\vec{s} \in \vec{S}$. Further, let $M \subseteq \vec{S}$ be some set of separations such that every $\overrightarrow{m} \in M$ is closely related to some profile Q_m of S which satisfies that $\vec{s} \in Q_m$. Then $\vec{r} := \vec{s} \wedge \bigwedge_{\overrightarrow{m} \in M} \overrightarrow{m}$ is an element of $\vec{S}$. Moreover, if $\vec{s}$ is closely related to some profile P of S, then $\vec{r}$ is closely related to P.*

The key tool in the proof of Theorem 1.1 is the following lemma which makes it possible to refine inessential stars that consist of separations whose inverses are each closely related to some $\mathcal{F}$-tangle. This lemma is taken from [2]; a less general version for submodular universes was first proved in [16].

LEMMA 3.4 ([2, Lemma 3.5]). — *Let $\vec{S}$ be a submodular separation system, and let $\mathcal{F}$ be a friendly set of stars in $\vec{S}$. Further, let $\sigma = \{\vec{s}_1, \dots, \vec{s}_k\} \subseteq \vec{S}$ be a star which is inessential for the set of all $\mathcal{F}$-tangles of S such that each $\overleftarrow{s}_i$ is closely related to some $\mathcal{F}$-tangle of S. Then there is an S-tree over $\mathcal{F} \cup \{\{\overleftarrow{s}_1\}, \dots, \{\overleftarrow{s}_k\}\}$ in which each $\vec{s}_i$ appears as a leaf separation.*

Erde [16] gave an example which shows that it is not possible to refine every inessential star, even if the separation system comes from a graph and $\mathcal{F} = \mathcal{T}_k^*$ for some $k \in \mathbb{N}$. Hence, it is not true that one can refine every nested set which distinguishes all the $\mathcal{F}$-tangles so that every inessential node of that refinement is a star in $\mathcal{F}$. For such a statement to be true, one has to impose further conditions on the nested sets.

A separation $s \in S$ is *good* (for a set $\mathcal{P}$ of profiles of S) if it distinguishes a pair of profiles $P, P' \in \mathcal{P}$ so that $\vec{s} \in P$ is closely related to P and $\overleftarrow{s} \in P'$ is closely related to P'. If we refer to a good set without reference to any set of profiles, then it can be assumed that this set is good for the set of all profiles of S.

We remark that "good" is a generalization of efficiency in that every separation in a submodular universe $\vec{U}$ which efficiently distinguishes some pair of profiles P, P' in U is good for $\{P, P'\}$ [2, Proposition 3.4]. Conversely, if a separation s of a graph G is good for a pair of regular profiles in G, then s distinguishes these two profiles efficiently [1, Example 5.1].

Lemma 3.4 implies that the inessential nodes of nested sets of good separations can be refined, which yields the following generalization of Theorem 1.1:

THEOREM 3.5 ([2, Theorem 2]). — *Let $\vec{S}$ be a submodular separation system, $\mathcal{F}$ be a friendly set of stars in $\vec{S}$, and let $\mathcal{P}$ be the set of all $\mathcal{F}$-tangles of S. Further, let $\widetilde{N} \subseteq S$ be a nested set that distinguishes all $\mathcal{F}$-tangles of S, and suppose that every separation in $\widetilde{N}$ is good for $\mathcal{P}$. Then there exists a nested set $N \subseteq S$ with $\widetilde{N} \subseteq N$ such that every node of N is either a star in $\mathcal{F}$ or home to an $\mathcal{F}$-tangle of S.*

4. Refining trees of tangles in graphs

Let G be a graph, $k \in \mathbb{N}$, and let $(\widetilde{T}, \widetilde{\mathcal{V}})$ be a tree-decomposition of G which distinguishes all its k-tangles so that every separation induced by an edge of $\widetilde{T}$ efficiently distinguishes some pair of k-tangles. Then Theorem 1.1 asserts that $(\widetilde{T}, \widetilde{\mathcal{V}})$ extends to a tree-decomposition all whose inessential parts are small. This is done by refining the inessential parts of $(\widetilde{T}, \widetilde{\mathcal{V}})$ with further tree-decompositions that decompose those parts into smaller parts. In this section we show that we can refine the essential parts with further tree-decompositions as well so that we then obtain a single tree-decomposition of G whose parts are all relatively small. The main result of this section is then Theorem 1.2, which extends Theorem 1.1 and also strengthens a result of Erde ([16, Theorem 4.6]).

In order to make the proof of Theorem 1.2 more straightforward, and to reduce clutter caused by translating back and forth between the tree-decompositions and the nested sets of separations they induce, we first prove Theorem 1.2′. Since every tree-decomposition of G induces a nested set of separations of G and vice versa, Theorem 1.2 then follows immediately.

THEOREM 1.2′. — *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a friendly set of stars in $\vec{S}_k(G)$. Further, let $\widetilde{N} \subseteq S_k(G)$ be a nested set that distinguishes all the $\mathcal{F}$-tangles of $S_k(G)$ so that every separation in $\widetilde{N}$ efficiently distinguishes a pair of $\mathcal{F}$-tangles of $S_k(G)$. Then there exist a nested set $N \subseteq S_k(G)$ with $\widetilde{N} \subseteq N$ such that*

- *(i) every inessential node of N is a star in $\mathcal{F}$;*
- *(ii) the interior of every essential node of N is of smallest size among all exclusive stars contained in the $\mathcal{F}$-tangle living at that node.*

Note that the set $\mathcal{T}_k^*$ of all stars in $\mathcal{T}_k$ satisfies the premise of Theorem 1.2′ (see Section 2), and hence we can apply Theorem 1.2′ to the set of all k-tangles in G for some $k \in \mathbb{N}$. In particular, we can use Theorem 1.2′ to refine the canonical nested set from [4], which efficiently distinguishes all the k-tangles in G:

COROLLARY 4.1. — *Let G be a graph, and $k \in \mathbb{N}$. Then there exist nested sets $\widetilde{N} \subseteq N \subseteq S_k(G)$ such that:*

- *(i) $\widetilde{N}$ is fixed under all automorphisms of G and efficiently distinguishes all k-tangles in G;*
- *(ii) every inessential node of N is a star in $\mathcal{T}_k$;*

(iii) the interior of every essential node of N is of smallest size among all exclusive stars contained in the k-tangle living at that node.

Before we prove Theorem 1.2′, let us first emphasize the optimality of its condition (ii). Since $\widetilde{N}$ distinguishes all $\mathcal{F}$-tangles of $S_k(G)$, the refinement N does so as well. Hence, each of its essential nodes is necessarily exclusive, and thus their interiors can only be of smallest size among all exclusive stars in the tangle they are home to. The following example shows that in general there may be non-exclusive stars in a tangle whose interiors contain fewer vertices than the interiors of the exclusive stars in that tangle. Thus, condition (ii) cannot be strengthened.

Example 4.2. — Let G be the graph depicted on the left side of Figure 4.1, which consists of five complete graphs K_{20} and one K_8 that are glued together as depicted on the left side of Figure 4.1. Whenever a K_{20} is denoted inside a face of the drawing, we assume that the K_{20} contains the boundary vertices and edges of this face, but no other depicted vertices or edges. Moreover, the K_8 contains all vertices in the grey cycle.

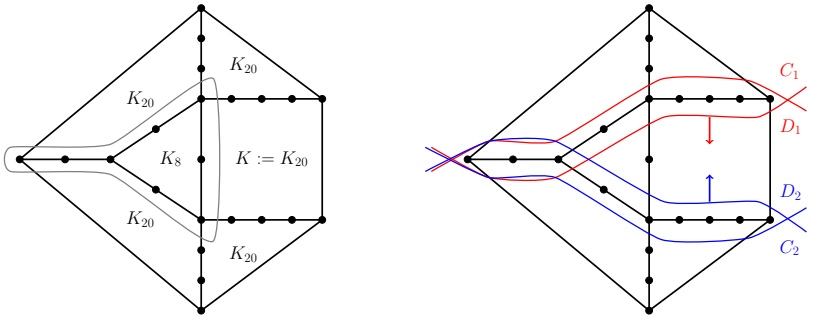

Figure 4.1. Example 4.2

Let τ be the 10-tangle defined by orienting every separation towards that side which contains the clique K. Further, observe that every separation $\{A, B\} \in S_{10}(G)$ with $V(K_8) \subseteq A \cap B$ is of the form $\{A, V(G)\}$. Thus, we obtain a consistent orientation τ' of $S_{10}(G)$ by letting $(A, V(G)) \in \tau'$ for every set $A \subseteq V(G)$ of at most nine vertices, and $(A, B) \in \tau'$ if $V(K_8) \subseteq B$ for all other separations $\{A, B\} \in S_{10}(G)$. It is straightforward to check that τ' is a 10-tangle in G.

Moreover, it is easy to see that every star ρ in τ whose interior is of smallest size among all stars in τ consists of $(C_1, D_1), (C_2, D_2)$ (see Figure 4.1 (right side)) and possibly some separations of the form $(A, V(G))$.

But these separations are also contained in τ', and thus ρ is not exclusive for the set of all 10-tangles in G.

Let us now turn to the proof of Theorem 1.2′. By Theorem 1.1, there exists a nested set $N' \subseteq S_k(G)$ which refines the inessential nodes of $\widetilde{N}$ and satisfies (i). Therefore, we are left to refine the essential nodes of N' so that this refinement satisfies (i) and (ii). For this, we show the following lemma about refining essential stars in graphs, which will imply Theorem 1.2′:

LEMMA 4.3. — *Let G be a graph, $k \in \mathbb{N}$, and let $\mathcal{F}$ be a friendly set of stars in $\vec{S}_k(G)$. Further, let $\sigma = \{\vec{s}_1, \dots, \vec{s}_n\} \subseteq \vec{S}_k(G)$ be a star which is home to a unique $\mathcal{F}$-tangle τ of $S_k(G)$, and suppose that each s_i efficiently distinguishes some pair of $\mathcal{F}$-tangles of $S_k(G)$.*

Then there exist a star $\sigma' \subseteq \tau$ whose interior is of smallest size among all exclusive stars in τ and an $S_k(G)$-tree over $\mathcal{F}' := \mathcal{F} \cup \{\sigma'\} \cup \{\{\overleftarrow{s}_1\}, \dots, \{\overleftarrow{s}_n\}\}$ in which each $\vec{s}_i$ appears as a leaf separation.

As τ has to live at a node of any $S_k(G)$-tree over $\mathcal{F}'$, and since σ' is the unique star in $\mathcal{F}'$ that is contained in τ, every $S_k(G)$-tree over $\mathcal{F}'$ has to contain a node which is associated with σ'. Thus, there can only exist an $S_k(G)$-tree over $\mathcal{F}'$ in which each $\vec{s}_i \in \sigma$ appears as a leaf separation if σ' is nested with σ. Our first step towards the proof of Lemma 4.3 will thus be to show that under all exclusive stars in τ whose interiors are of smallest size there is at least one star which is nested with σ.

LEMMA 4.4. — *Let $k \in \mathbb{N}$, let $\mathcal{P}$ be some set of k-profiles in a graph G, and let $P \in \mathcal{P}$. Further, let $\sigma \subseteq P$ be a star, and suppose that every separation in σ efficiently distinguishes some pair of k-profiles in $\mathcal{P}$. Then there exists a star $\rho \subseteq P$ with $\sigma \leqslant \rho$ whose interior is of smallest size among all stars in P that are exclusive for $\mathcal{P}$.*

Proof. — Let ρ be a star in P which is exclusive for $\mathcal{P}$ and which is nested with as many separations in σ as possible such that its interior is of smallest size among all stars in P that are exclusive for $\mathcal{P}$. We claim that ρ is in fact nested with σ. This then clearly implies the assertion. Indeed, the unique node ρ' of $\{\{A, B\} : (A, B) \in \sigma \cup \rho\}$ which is home to P then satisfies both $\sigma \leqslant \rho'$ and $\rho \leqslant \rho'$, and in particular $\mathrm{int}(\rho') \subseteq \mathrm{int}(\rho)$. Hence, ρ' is as desired.

So suppose for a contradiction that there is a separation $(C, D) \in \sigma$ which is not nested with ρ. By assumption, $\{C, D\}$ efficiently distinguishes some pair $Q \ni (D, C)$ and Q' of k-profiles in $\mathcal{P}$. Since Q is consistent and because $\rho \not\subseteq Q$ as ρ is P-exclusive for $\mathcal{P}$, there is a unique separation

$(A,B) \in \rho$ with $(B,A) \in Q$. Then $(C,D) \wedge (A,B)$ and $(C,D) \wedge (B',A')$, for every $(A',B') \in \rho \setminus \{(A,B)\}$, have order at least $|C \cap D|$ because $\{C,D\}$ distinguishes Q and Q' efficiently. By submodularity, it follows that $(C,D) \vee (B',A')$ has order $\leqslant |A' \cap B'| < k$, and further

$$\begin{aligned} |A \cap B| &\geqslant |(C \cup A) \cap (D \cup B)| = |(A \cap B \cap D) \cup (B \cap C \cap D)| \\ &= |A \cap B \cap (D \setminus C)| + |B \cap C \cap D|. \end{aligned} \tag{4.1}$$

Thus, the star $\rho' := \{(A \cup C, B \cap D)\} \cup \{(A' \cap D, B' \cup C) : (A',B') \in \rho \setminus \{(A,B)\}\}$ is a subset of $\vec{S}_k(G)$. In particular, since P is a profile, we have $\rho' \subseteq P$. Moreover, ρ' is P-exclusive for $\mathcal{P}$ since for every k-profile $P' \neq P$ in $\mathcal{P}$ either $(B,A) \in P'$ or $(D,C) \in P'$ and thus $(B \cap D, A \cup C) \in P'$, or there exists some $(A',B') \in \rho$ such that $(B',A') \in P'$, and we then have $(B' \cup C, A' \cap D) \in P'$ because P' is a profile.

We claim that the interior X' of ρ' has at most the size of the interior X of ρ. This then contradicts the choice of ρ since ρ' is nested with $(C,D) \in \sigma$ by construction, and thus, by Lemma 2.1, ρ' is nested with at least one separation in σ more than ρ. Indeed, we have

$$\begin{aligned} X' &= (B \cap D) \cap \bigcap_{(A',B') \in \rho \setminus \{(A,B)\}} (B' \cup C) = (X \cap D) \cup (B \cap C \cap D) \\ &= (X \cap (D \setminus C)) \,\dot\cup\, (B \cap C \cap D). \end{aligned}$$

Since $\{C,D\}$ is a separation of G, we have $X = (X \cap C) \,\dot\cup\, (X \cap (D \setminus C))$, and thus

$$|X'| = |X| - |X \cap C| + |B \cap C \cap D| \leqslant |X| - |(A \cap B) \cap C| + |B \cap C \cap D|,$$

where we have used that $A \cap B \subseteq X$. Combining this with (4.1) yields that

$$|X'| \leqslant |X| - |(A \cap B) \cap C| + (|A \cap B| - |A \cap B \cap (D \setminus C)|) = |X|$$

since $|A \cap B| = |(A \cap B) \cap C| + |A \cap B \cap (D \setminus C)|$, which concludes the proof. □

For the proof of Lemma 4.3 we need one further auxiliary statement, Lemma 4.7 below, which follows from the following two lemmas that are immediate from [16, Lemmas 4.2 & 4.3].

Lemma 4.5. — *Let $\vec{s} \leqslant \vec{r}$ be two separations of some graph G. If $\vec{x}$ is a separation of G of minimal order such that $\vec{s} \leqslant \vec{t} \leqslant \vec{r}$, then $|\vec{t} \wedge \vec{x}| \leqslant |\vec{x}|$ for all $\vec{x} \leqslant \vec{r}$.*

Lemma 4.6. — *Let $\vec{r}_1$ and $\vec{r}_2$ be two separations of some graph G. Let $\vec{s}_1$ be any separation of minimal order such that $\vec{r}_1 \wedge \overleftarrow{r}_2 \leqslant \vec{s}_1 \leqslant \vec{r}_1$ and set $\vec{s}_2 := \vec{r}_2 \wedge \overleftarrow{s}_1$. Then $|\vec{s}_i| \leqslant |\vec{r}_i|$ for $i \in \{1,2\}$, and for $\vec{x} \leqslant \vec{r}_i$ we have $|\vec{s}_i \wedge \vec{x}| \leqslant |\vec{x}|$.*

LEMMA 4.7. — *Let P be a k-profile in some graph G for some $k \in \mathbb{N}$, and let $\vec{s} \in P$. Further, let $\overrightarrow{s'} \in P$ be any separation of minimal order such that $\vec{s} \leqslant \overrightarrow{s'}$. Then $\overrightarrow{s'}$ is closely related to P, and for every $\vec{r} \in P$ with $\vec{r} \leqslant \overleftarrow{s}$ we have $\vec{r} \wedge \overleftarrow{s'} \in \vec{S}_k(G)$. Moreover, if $\vec{r}$ is closely related to P, then $\vec{r} \wedge \overleftarrow{s'}$ is closely related to P.*

Proof. — Let $\vec{t} \in P$ be a maximal separation in P such that $\overrightarrow{s'} \leqslant \vec{t}$. Applying Lemma 4.5 to $\vec{s}$ and $\vec{t}$ yields that $\overrightarrow{s'} \wedge \vec{x} \in \vec{S}_k(G)$ for every $\vec{x} \leqslant \vec{t}$, which by Propositions 3.1 and 3.2 implies that $\overrightarrow{s'}$ is closely related to P. Further, applying Lemma 4.6 to $\overrightarrow{s'}$ and $\vec{r}$ yields that $\vec{r} \wedge \overleftarrow{s'}$ is an element of $\vec{S}_k(G)$ and that $(\vec{r} \wedge \overleftarrow{s'}) \wedge \vec{x} \in \vec{S}_k(G)$ for every $\vec{x} \leqslant \vec{r}$, which by Proposition 3.2 implies that $\vec{r} \wedge \overleftarrow{s'}$ is closely related to P if $\vec{r}$ is closely related to P. □

We can now prove Lemma 4.3:

Proof of Lemma 4.3. — Let $\mathcal{P}$ be the set of all $\mathcal{F}$-tangles of $S_k(G)$, and first assume that there is a star $\sigma' \subseteq \tau$ with $\sigma \leqslant \sigma'$ whose interior is of smallest size among all stars in τ that are exclusive for $\mathcal{P}$, and which has the further property that every separation in σ' is closely related to τ. Then the assertions follows. Indeed, set $N' := \{s : \vec{s} \in \sigma \cup \sigma'\}$ and let ρ be an inessential node of N'. Then $\rho \neq \sigma'$ since σ' is home to τ and hence essential. In particular, ρ contains no separation from σ'. Thus, by the assumptions on σ and σ', every separation in ρ has an inverse that is closely related to an $\mathcal{F}$-tangle of $S_k(G)$, which implies that ρ satisfies the premise of Lemma 3.4. So we may apply Lemma 3.4 to the inessential nodes of N', which yields a refinement N of N' all whose inessential nodes are stars in $\mathcal{F}$. This set N induces an $S_k(G)$-tree over $\mathcal{F} \cup \{\sigma'\} \cup \{\{\overleftarrow{s}_1\}, \dots, \{\overleftarrow{s}_n\}\}$ in which each $\vec{s}_i$ appears as a leaf separation. By the choice of σ', this clearly implies the assertion.

Thus, it suffices to find a star σ' as above. For this, let $\sigma' \subseteq \tau$ be a star such that

(1) $\sigma \leqslant \sigma'$ and $|\mathrm{int}(\sigma')| = \min\{|\mathrm{int}(\rho)| : \rho \subseteq \tau \text{ and } \rho \text{ is exclusive for } \mathcal{P}\}$, and
(2) the number of separations in σ' that are not closely related to τ is minimal among all stars in τ that satisfy (1).

Note that, by Lemma 4.4, there exists a star which satisfies (1). We claim that σ' is closely related to τ, which implies that σ' is the star for the argument above, and thus concludes the proof. For this, suppose for a contradiction that there is a separation $(C', D') \in \sigma'$ which is not closely related to τ, and let $(A, B) \in \tau$ be of minimal order such that $(C', D') \leqslant$

(A, B). We first show that there is such a separation $\{A, B\}$ which is nested with σ.

Claim 4.8. — *For every $\vec{r} \in \tau$ there exists a separation $\overrightarrow{r'} \in \tau$ of order $\leqslant |r|$ such that r' is nested with σ. Moreover, if $\vec{s} \in \tau$ is nested with σ such that $\vec{s} \leqslant \vec{r}$, then r' can be chosen so that $\vec{s} \leqslant \overrightarrow{r'}$.*

Proof of the Claim. — Let $\overrightarrow{r'}(\geqslant \vec{s})$ be a separation of order $\leqslant |r|$ which is nested with as many separations in σ as possible. We claim that $\overrightarrow{r'}$ is nested with σ, which then clearly implies the assertion. For this, suppose for a contradiction that $\overrightarrow{r'}$ crosses some $\vec{t} \in \sigma$. By the assumption on σ, there exists a pair of k-profiles $Q \ni \overleftarrow{t}$ and $Q' \ni \vec{t}$ in G such that t distinguishes them efficiently. Since Q is a k-profile, it contains an orientation of r'; by symmetry we may assume that $\overrightarrow{r'} \in Q$. If $\overrightarrow{r'} \vee \overleftarrow{t} \in \vec{S}_k(G)$, then $\overrightarrow{r'} \vee \overleftarrow{t} \in Q$ and $\overleftarrow{r'} \wedge \vec{t} \in Q'$ as Q and Q' are k-profiles, and thus $\vec{r} \vee \overleftarrow{t}$ distinguishes Q and Q'. Since t efficiently distinguishes Q and Q', this implies that $\overrightarrow{r'} \vee \overleftarrow{t}$ has order $\geqslant |t|$, and hence, by submodularity, $\overrightarrow{r'} \wedge \overleftarrow{t}$ has order $\leqslant |r'|$. Then $\overrightarrow{r'} \wedge \overleftarrow{t} \in \tau$ because τ is a profile, which contradicts the choice of $\overrightarrow{r'}$ since $\overrightarrow{r'} \wedge \overleftarrow{t}$ is nested with one separation in σ more than $\overrightarrow{r'}$ by Lemma 2.1 (and still $\vec{s} \leqslant \overrightarrow{r'} \wedge \overleftarrow{t}$ as $\vec{s} \leqslant \vec{r}$ and s is nested with t). □

By Claim 4.8, we may assume that $\{A, B\}$ is nested with σ. Then, by applying Lemma 4.7 to $\vec{s} := (C', D')$ and $\overrightarrow{s'} := (A, B)$, we obtain that the star

$$\sigma'' := \{(A, B)\} \cup \{(C \cap B, D \cup A) : (C, D) \in \sigma' \setminus \{(C', D')\}\}$$

is a subset of $\vec{S}_k(G)$, which by the consistency of τ implies that $\sigma'' \subseteq \tau$. Moreover, we have $\sigma \leqslant \sigma''$ as $\sigma \leqslant \sigma'$ and (A, B) is nested with σ. We claim that σ'' witnesses that σ' does not satisfy (2), which then contradicts the choice of σ' and thus concludes the proof. By Lemma 4.7, (A, B) is closely related to τ, and if $(C, D) \in \sigma'$ is closely related to τ, then $(C \cap B, D \cup A)$ is closely related to τ, too. Thus, there are fewer separations in σ'' that are not closely related to τ than in σ'. Therefore, we are left to show that σ'' satisfies (1). By the argument above, we have $\sigma \leqslant \sigma''$, so we only need to show that the interior X'' of σ'' has at most the size of the interior X' of σ'. By definition, we have

$$\begin{aligned} X'' &= B \cap \bigcap_{(C,D)\in\sigma'} (D \cup A) = B \cap (X' \cup A) \\ &= (B \cap X') \cup (B \cap A) = (B \cap X') \,\dot\cup\, ((B \cap A) \setminus X'). \end{aligned}$$

Since $\{A, B\}$ is a separation of G, we have $X' = (X' \cap B) \dot\cup (X' \cap (A \setminus B))$ and thus

$$|X''| = |X'| - |(A \setminus B) \cap X'| + |(A \cap B) \setminus X'|.$$

So we are done if $|(A \cap B) \setminus X'| \leqslant |(A \setminus B) \cap X'|$. Set $\rho := \sigma' \setminus \{(C', D')\}$. By the choice of (A, B), the separation $(A, B) \wedge \bigwedge_{(C,D)\in\rho}(D, C)$ has at least order $|A \cap B|$, and thus

$$\begin{aligned} |A \cap B| &\leqslant \Bigg| A \cap \bigcap_{(C,D)\in\rho} D \cap \Bigg(B \cup \bigcup_{(C,D)\in\rho} C \Bigg) \Bigg| \\ &= \Bigg| \Bigg(A \cap B \cap \bigcap_{(C,D)\in\rho} D \Bigg) \cup \Bigg(A \cap \bigcap_{(C,D)\in\rho} D \cap \bigcup_{(C,D)\in\rho} C \Bigg) \Bigg|. \end{aligned}$$

Since σ' is a star, we have $C \subseteq D'$ for every $(C, D) \in \rho$, and hence $\bigcup_{(C,D)\in\rho} C \subseteq D'$. Moreover, by the choice of (A, B), we have $B \subseteq D'$, and thus

$$|A \cap B| \leqslant |(A \cap B \cap X') \cup (A \cap X')| = |A \cap B \cap X'| + |(A \setminus B) \cap X'|.$$

Combining this inequality with $|A \cap B| = |A \cap B \cap X'| + |(A \cap B) \setminus X'|$ yields that

$$\begin{aligned} |(A \setminus B) \cap X'| &\geqslant |A \cap B| - |A \cap B \cap X'| \\ &= |A \cap B| - (|A \cap B| - |(A \cap B) \setminus X'|) = |(A \cap B) \setminus X'|, \end{aligned}$$

which concludes the proof. □

With Lemma 4.3 at hand Theorem 1.2′ follows immediately:

Proof of Theorem 1.2′. — Apply Theorem 1.1 to $\widetilde{N}$ to obtain a nested set $N' \subseteq S_k(G)$ with $\widetilde{N} \subseteq N'$ which satisfies (i). Then, apply Lemma 4.3 to the essential nodes of N'. □

Proof of Theorem 1.2. — Apply Theorem 1.2′ to the set of separations induced by $(\widetilde{T}, \widetilde{\mathcal{V}})$. □

Recall that the essential parts of the tree-decomposition from Theorem 1.2 are as small as possible so that they are still home to their tangle. Thus, this refinement is optimal in that it decreases the size of the essential parts as much as possible.

Now, one may ask what else could be said about the essential parts of that refinement besides their size being as small as possible. Our main reason for decreasing the size of the parts was to more precisely exhibit where the tangles are located in the graph, so it seems natural to ask whether the vertices in the essential parts in some sense "belong" to the tangle living in

that part, and additionally whether all vertices that "belong" to a tangle are contained in the part which is home to that tangle.

In general, this question cannot be answered because tangles are, by nature, a somewhat fuzzy object, which is, at least in general, not related to a specific set of vertices. So usually it cannot be determined at all whether a certain vertex "belongs" to a tangle or not. Thus, in general, we also cannot evaluate whether the essential parts contain precisely those vertices which "belong" to that tangle.

However, in the following, we present two properties of the refined tree-decomposition from Theorem 1.2 that we believe validate that its essential parts, at least in many cases, "closely correspond" to the tangle they are home to. For this, let G be some graph, and let $(T, \mathcal{V})$ be the tree-decomposition which one obtains by applying Theorem 1.2 to any tree-decomposition of G which efficiently distinguishes all the regular k-profiles in G. (Such tree-decompositions exist by [13, Theorem 3.6].)

Then each part of $(T, \mathcal{V})$ of size $> 3k - 3$ is home to a k-tangle, which it then "witnesses" and "induces". More precisely, let us say that a subgraph H of G *witnesses* a k-tangle τ if for every subset $\{(A_i, B_i) : i \in [3]\}$ of τ of at most three elements we have $\bigcup_{i\in[3]} G[A_i \cap V(H)] \neq H$. Further, a set U of vertices of G *induces* a k-profile P if for every separation $(A, B) \in P$ we have $|A \cap U| < |B \cap U|$.

If a part $V_t \in \mathcal{V}$ has size $> 3k-3$, then σ_t cannot be a star in $\mathcal{P}_{S_k(G)} \subseteq \mathcal{T}_k$, and thus by (i), V_t is home to some regular k-profile τ in G. By (ii) and the following lemma, it then follows that $|A \cap V_t| < k$ for every $(A, B) \in \tau$.

Lemma 4.9 ([11, Proof of Theorem 12]). — *Let $k \in \mathbb{N}$, let $\mathcal{P}$ be some set of k-profiles in a graph G, and let $P \in \mathcal{P}$. Further, let $\sigma \subseteq P$ be a star whose interior is of smallest size among all stars in P that are exclusive for $\mathcal{P}$. Then $|A \cap \operatorname{int}(\sigma)| < k$ for all $(A, B) \in P$.*

Note that in [11] the statement of Lemma 4.9 is shown for stars in P whose interior is of smallest size among all stars in P. However, it is easy to see that the same proof also yields the statement of Lemma 4.9, where the interior of σ is only of smallest size among all *exclusive* stars in P.

By Lemma 4.9 we have $|A \cap V_t| < k \leqslant (3k-2)-(k-1) \leqslant |V_t|-|A \cap V_t| = |(B \setminus A) \cap V_t| \leqslant |B \cap V_t|$, and hence V_t induces τ. Moreover, for every subset $\{(A_i, B_i) : i \in [3]\} \subseteq \tau$, we have $\bigcup_{i\in[3]} G[A_i \cap V_t] \neq G[V_t]$ since $|V_t \cap \bigcap_{i\in[3]}(B_i \setminus A_i)| \geqslant |V_t| - |A_1| - |A_2| - |A_3| \geqslant (3k-2) - 3(k-1) = 1$. Hence, τ is in fact a k-tangle in G as witnessed by $G[V_t]$.

Therefore, every large enough part of $(T, \mathcal{V})$ induces a k-tangle. Conversely, if a k-tangle is induced by some large and highly connected substructure of G, does that substructure appear as a part of $(T, \mathcal{V})$? In general, this will not be the case even if a k-tangle is induced by a large clique, since even a clique must not be equal to a part of any tree-decomposition of adhesion $< k$ [5]. However, if a k-tangle is induced by a "k-block", then the k-block will be equal to a part of $(T, \mathcal{V})$ if there exists any tree-decomposition of adhesion $< k$ at all that contains the k-block as a part.

For some $k \in \mathbb{N}$, a *k-block* in a graph G is a maximal set b of at least k vertices such that no two vertices $v, w \in b$ can be separated in G by fewer than k vertices. It is straight forward to check that every k-block *induces* a regular k-profile by orienting $\{A, B\} \in S_k(G)$ as (A, B) if and only if $b \subseteq B$. A k-block b in G is *separable* if it is the interior of some star in $S_k(G)$, i.e. if there exists a star $\sigma \subseteq \vec{S}_k(G)$ such that $\mathrm{int}(\sigma) = b$.

Now suppose that b is a separable k-block in G, and let τ be the k-profile induced by b. Then the star ρ that consists precisely of the separations $(V(C) \cup N(C), V(G) \setminus V(C))$ for components C of $G - b$ is in $\vec{S}_k(G)$ [6, Lemma 4.1], and it is easy to see that $\mathrm{int}(\rho) = b$ and that τ is the unique k-profile which contains ρ. Moreover, we have $b \subseteq B$ for every $(A, B) \in \tau$, and thus $b \subseteq \mathrm{int}(\sigma)$ for every star $\sigma \subseteq \tau$. Hence, if b is separable, then every star in τ whose interior is of smallest size among all exclusive stars in τ will be equal to b. It follows by condition (ii) that the essential part of the tree-decomposition $(T, \mathcal{V})$ from Theorem 1.2 which is home to τ is equal to b. Thus, every separable k-block in G appears as a part of $(T, \mathcal{V})$. All in all, we have the following theorem:

THEOREM 4.10. — *Let G be a graph and $k \in \mathbb{N}$. Then there exists a tree-decomposition $(T, \mathcal{V})$ of G of adhesion $< k$ which has the following properties:*

(i) $(T, \mathcal{V})$ efficiently distinguishes all regular k-profiles in G;
(ii) each part of $(T, \mathcal{V})$ of size $> 3k - 3$ is home to a k-tangle, which it then witnesses and induces;
(iii) every separable k-block in G appears as a part of $(T, \mathcal{V})$.

We remark that the existence of tree-decompositions satisfying (i) and (iii) is already known to Carmesin and Gollin [6] (see also [3] for a short proof). What is new is property (ii): that all the inessential parts of the tree-decomposition $(T, \mathcal{V})$ are small, while all of its parts that are not small induce k-tangles.

5. Refining trees of tangles in abstract separation systems

In this section we extend Theorem 1.2 to abstract separation systems. For this note that condition (i) of Theorem 1.2′ is already formulated in a way which translates verbatim to abstract separation systems, while its condition (ii) makes use of the specific properties of graphs. More precisely, (ii) refers to the size of the interiors of the essential nodes, i.e. the number of vertices contained in them. Since vertices need not exist in the more general setting of abstract separation systems, we need to adapt this condition in order to generalize Theorem 1.2.

For this, recall that the purpose of condition (ii) is to ensure that the refinement is optimal. More precisely, (ii) makes sure that the interiors of the essential nodes of the refined nested set are as small as possible and thereby guarantees that the essential nodes are as "close" to the tangle they are home to as possible. Here "small" refers to the number of vertices in the interiors. Since we cannot rely on the existence of vertices, as mentioned above, we need to find a different measure to determine the "closeness" of the essential nodes to the tangle they are home to. In particular, as the notion of an "interior" of a node does not extend to abstract separation systems, this measure needs to refer to the nodes themselves.

Since the nodes of a nested set are stars, we may consider the partial order on them. For this, recall that if S is some separation system, then the partial order on $\vec{S}$ induces a partial order on the set of all proper stars in $\vec{S}$ (see Section 2.1). Clearly, for a given profile P of S, those stars in P that are maximal among all stars in P are closest to P. We thus replace (ii) in Theorem 1.2′ with the condition that every essential node is a maximal star in the tangle living at it. We then obtain Theorem 1.3, which is the main result of this section, and which we restate here. For this, set $\mathcal{T}' := \{\{\vec{r}, \vec{s}\} \subseteq \vec{S} : \vec{r} \leqslant \overleftarrow{s} \text{ and } \vec{r} \vee \vec{s} \text{ is co-small}\}$.

Theorem 1.3. — *Let $\vec{S}$ be a submodular separation system inside some distributive universe, and let $\mathcal{F}$ be a friendly set of stars in $\vec{S}$ with $\mathcal{T}' \subseteq \mathcal{F}$. Further, let $\widetilde{N} \subseteq S$ be a nested set that distinguishes all $\mathcal{F}$-tangles of S, and suppose that every separation in $\widetilde{N}$ is good for the set of all $\mathcal{F}$-tangles of S. Then there is a nested set $N \subseteq S$ with $\widetilde{N} \subseteq N$ such that*

(i) every inessential node of N is an element of $\mathcal{F}$;

(ii) every essential node of N is maximal in the $\mathcal{F}$-tangle living at it.

Note that by [2, Theorem 5.8] there exists a nested set $\widetilde{N}$ of separations which satisfies the premise of Theorem 1.3. This set is even canonical, but in a slightly weaker sense than usual.

Further, recall that "good" generalizes the notion of efficiency to separations systems that come without an order function (see Section 3).

Before we prove Theorem 1.3, let us first make some remarks about it. The set N refines the nodes of $\widetilde{N}$ in that it "splits" them into smaller nodes which are each either a star in $\mathcal{F}$ if they are inessential, or "close" to the tangle they are home to. More precisely, (i) implies that every inessential node of the refinement is in fact too small to be home to a tangle, while (ii) guarantees that the essential nodes are as "close" to the tangle they are home to as possible, where the "closeness" is measured by the partial order on the set of stars. In particular, condition (ii) ensures that the essential nodes of the refinement cannot be refined any further, and hence the refinement from Theorem 1.3 is optimal in that sense.

An *abstract tangle* of a separation system S is a $\widetilde{\mathcal{T}}$-tangle where $\widetilde{\mathcal{T}} := \big\{\{\vec{r},\vec{s},\vec{t}\} \subseteq \vec{S} : \vec{r} \vee \vec{s} \vee \vec{t} \text{ is co-small}\big\}$ [12]. Since $\widetilde{\mathcal{T}}^*$ (the set of all stars in $\widetilde{\mathcal{T}}$) is closed under shifting and $\widetilde{\mathcal{T}}^* \supseteq \mathcal{T}'$, we can apply Theorem 1.3 to the set of all abstract tangles if S lies inside some distributive universe.

Moreover, since every separation system of the form S_k inside a submodular universe is submodular, Theorem 1.3 can also be applied to such systems. In fact, we obtain the following corollary by applying it to the nested set which Diestel, Hundertmark and Lemancyzk [13, Theorem 3.6] constructed:

Corollary 5.1. — *Let $\vec{U}$ be a distributive, submodular universe of separations, and let $k \in \mathbb{N}$. Further, let $\mathcal{F}$ be a friendly set of stars in $\vec{S}_k$ such that $\mathcal{T}' \subseteq \mathcal{F}$, and let $\mathcal{P}$ be the set of all $\mathcal{F}$-tangles of S_k. Then there exist nested sets $\widetilde{N} := \widetilde{N}(\mathcal{P}) \subseteq N \subseteq S_k$ such that*

(i) $\widetilde{N}$ efficiently distinguishes all tangles in $\mathcal{P}$;
(ii) for every automorphism α of $\vec{U}$ we have $\widetilde{N}(\mathcal{P}^\alpha) = \widetilde{N}(\mathcal{P})^\alpha$;
(iii) every inessential node in N is a star in $\mathcal{F}$;
(iv) every essential node in N is maximal in the $\mathcal{F}$-tangle living at it.

Proof. — By [13, Theorem 3.6] there exists a nested set $\widetilde{N} \subseteq S_k$ which satisfies (i) and (ii) such that every separation in $\widetilde{N}$ distinguishes some pair of profiles in $\mathcal{P}$ efficiently. Since every separation which efficiently distinguishes some pair of profiles in $\mathcal{P}$ is good for $\mathcal{P}$ by [2, Proposition 3.4], applying Theorem 1.3 to $\widetilde{N}$ yields the existence of a nested set $N \subseteq S_k$ with $\widetilde{N} \subseteq N$ that satisfies (iii) and (iv). □

Let us now turn to the proof of Theorem 1.3. For this, let a submodular separation system $\vec{S}$, a set $\mathcal{F}$ of stars in $\vec{S}$ and a nested set $\widetilde{N} \subseteq S$ be given which satisfy the premise of Theorem 1.3. By Theorem 3.5, there exists a

refinement N' of $\widetilde{N}$ which satisfies (i), i.e. which refines all the inessential nodes of $\widetilde{N}$ so that all inessential nodes are stars in $\mathcal{F}$. Therefore, in order to prove Theorem 1.3, we only need to show that we can refine the essential nodes too, i.e. that we can "split" the essential nodes into smaller nodes so that all inessential nodes are stars in $\mathcal{F}$, and the new essential nodes are maximal in the tangle they are home to. For this, we prove the following variant of Lemma 4.3:

LEMMA 5.2. — *Let $\vec{S}$ be a submodular separation system inside a distributive universe $\vec{U}$, and let $\mathcal{F} \supseteq \mathcal{T}'$ be a friendly set of stars in $\vec{S}$. Further, let $\sigma = \{\vec{s}_1, \dots, \vec{s}_k\} \subseteq \vec{S}$ be a star which is home to a unique $\mathcal{F}$-tangle P of S, and suppose that each s_i is good for the set of all $\mathcal{F}$-tangles of S. Then there exist a maximal star $\sigma' \subseteq P$ and an S-tree over $\mathcal{F}' := \mathcal{F} \cup \{\sigma'\} \cup \{\{\overline{s}_1\}, \dots, \{\overline{s}_k\}\}$ in which each $\vec{s}_i$ appears as a leaf separation.*

Just as Theorem 1.1 and Lemma 4.3 imply Theorem 1.2, their generalizations to abstract separation systems (Theorem 3.5 and Lemma 5.2) imply Theorem 1.3.

Proof of Theorem 1.3. — By applying Theorem 3.5 to $\widetilde{N}$ we obtain a refinement N' of $\widetilde{N}$ whose inessential nodes are all stars in $\mathcal{F}$. Let $\mathcal{P}$ be the set of all $\mathcal{F}$-tangles of S, let $P \in \mathcal{P}$ be given, and let σ be the node of N' that is home to P. By Lemma 5.2, there exists a star $\sigma' \subseteq P$ with $\sigma \leqslant \sigma'$ that is maximal in P, and an S-tree over $\mathcal{F} \cup \{\sigma'\} \cup \{\{\overline{s}\} : \vec{s} \in \sigma\}$ in which each $\vec{s} \in \sigma$ appears as a leaf separation. Let N_P be the set of separations induced by that S-tree. Clearly, N_P is nested with N'. Setting $N := N' \cup \bigcup_{P \in \mathcal{P}} N_P$ then yields the claim. □

Let $\sigma \subseteq \vec{S}$ be a star that satisfies the premise of Lemma 5.2, and let P be the unique $\mathcal{F}$-tangle living at σ. One particularly natural approach to construct an S-tree as in Lemma 5.2 is the following. Take a maximal star $\sigma' \subseteq P$ with $\sigma \leqslant \sigma'$, and consider the set $N'_P := \{s : \vec{s} \in \sigma \cup \sigma'\}$. If σ is not already maximal in P, then $\sigma < \sigma'$, which implies that N'_P has some inessential nodes. As these nodes will in general not be stars in $\mathcal{F}$, we need to refine them further with S-trees as in Lemma 3.4. If that is possible, i.e. if there exists for every inessential node ρ of N'_P an S-tree over $\mathcal{F} \cup \{\{\overline{r}\} : \vec{r} \in \rho\}$, then these S-trees give rise to one overall S-tree as in Lemma 5.2.

However, this seemingly easy approach entails some difficulties. The refining lemma for inessential stars (Lemma 3.4) relies on the assumption that the inverse of every separation in the inessential star which we want

to refine is closely related to some $\mathcal{F}$-tangle of S. Since this assumption is necessary, as pointed out earlier, we cannot pick the star σ' arbitrarily among the maximal stars in P. Indeed, as the inverse of every separation in $\sigma' \setminus \sigma$ will be contained in an inessential node of N'_P, we have to pick a star $\sigma' \subseteq P$ which is closely related to P in order to be able to refine the newly arising inessential nodes. Unfortunately though, we do not know whether there always exists a star which is maximal in P and closely related to P.

Still, we hold on to the strategy of first choosing a suitable star σ' in P with $\sigma \leqslant \sigma'$, and then refining the inessential nodes of N'_P. To make this strategy work, we relax the conditions that we impose on the star σ'. Instead of requiring σ' to be a maximal star in P, we only ask σ' to be "near-maximal" in P, a slightly weaker property that will allow us to find a star σ' which is "near-maximal" in and closely related to P. Let us first make this property formal.

We call a set $R \subseteq P$ *narrow in* P if for every $\vec{x} \in P$ the separation $\overleftarrow{x} \vee \bigvee_{\vec{r} \in R} \vec{r}$ is co-small. Further, a star $\sigma \subseteq P$ is *near-maximal in* P if σ is narrow in P and for every $\vec{x} \in P$ there is at most one $\vec{s} \in \sigma$ such that $\vec{s} \leqslant \vec{x}$.

The main part of the proof of Lemma 5.2 is to show that we can always find a star $\sigma' \geqslant \sigma$ which is both closely related to and near-maximal in P.

LEMMA 5.3. — *Let $\vec{S}$ be a submodular separation system inside some distributive universe, and let P be a profile of S. Further, let $\sigma \subseteq P$ be a star which is good for the set of all profiles of S. Then there exists a star $\sigma' \subseteq P$ with $\sigma \leqslant \sigma'$ which is closely related to and near-maximal in P.*

With Lemma 5.3 at hand, we can show Lemma 5.2:

Proof of Lemma 5.2. — By applying Lemma 5.3 to σ we obtain a star $\sigma' \subseteq P$ with $\sigma \leqslant \sigma'$ which is both closely related to and near-maximal in P. Set $N' := \{s : \vec{s} \in \sigma \cup \sigma'\}$, and let Σ be the set of inessential nodes of N'. By assumption, every inessential node $\rho \in \Sigma$ contains only separations whose inverse is closely related to some $\mathcal{F}$-tangle of S. Indeed, since σ' is home to P and hence essential, we have $\rho \neq \sigma'$. In particular, ρ contains no separation from σ'. Thus, by the assumptions on σ and σ', the star ρ satisfies the premise of Lemma 3.4. So we may apply Lemma 3.4 to ρ to obtain an S-tree over $\mathcal{F} \cup \{\{\overleftarrow{r}\} : \vec{r} \in \rho\}$ in which each $\vec{r} \in \rho$ appears as a leaf separation. Let N_ρ be the set of separations induced by that S-tree. Clearly, N_ρ is nested with N'.

Take a maximal star $\sigma'' \subseteq P$ with $\sigma' \leqslant \sigma''$. By the definition of near-maximal, every inessential node ρ of $\{s : \vec{s} \in \sigma' \cup \sigma''\}$ has one of the following two forms. Either $\rho = \{\overleftarrow{s}\}$ for a separation $\vec{s} \in \sigma''$, which then must be small, or $\rho = \{\vec{r}, \overleftarrow{s}\}$ where $\vec{r} \in \sigma'$ and $\vec{s} \in \sigma''$ such that $\vec{r} \vee \overleftarrow{s}$ is

co-small. In the first case, we have $\rho \in \mathcal{F}$ because $\mathcal{F}$ is friendly, and in the second case we have $\rho \in \mathcal{T}' \subseteq \mathcal{F}$.

Thus, $\widetilde{N} := N' \cup \bigcup\{N_\rho : \rho \in \Sigma\} \cup \{s : \vec{s} \in \sigma''\}$ is a nested set of separations that induces an S-tree over $\mathcal{F} \cup \{\sigma''\} \cup \{\{\overleftarrow{s}\} : \vec{s} \in \sigma\}$ in which each $\vec{s} \in \sigma$ appears as a leaf separation. □

The remainder of this section is devoted to the proof of Lemma 5.3. For this, for the remainder of this section, let S be a submodular separation system. Further, let P be a profile of S, and let $\sigma \subseteq P$ be a *star of good separations*, that is a star of separations which are all good for the set of all profiles of S. We define $P_\sigma := \{\vec{s} \in P : \vec{s} \text{ is nested with } \sigma\}$ to be the set of all separations in P that are nested with σ.

The proof of Lemma 5.3 consists of two steps. First, we show that there exists a star in P_σ which is both closely related to and narrow in P. Then, we prove that every star in P_σ which contains the least number of separations such that it is closely related to and narrow in P is near-maximal in P. This then concludes the proof of Lemma 5.3.

To find a star $\sigma' \subseteq P_\sigma$ which is both closely related to and narrow in P, we consider the set R^P_σ of all maximal separations in P_σ. By Proposition 3.1, R^P_σ is closely related to P. Moreover, it is easy to check, as we will do later, that R^P_σ is narrow in P. However, R^P_σ is in general not nested and hence not a star. Therefore, in the following, we introduce a method[(1)] of turning R^P_σ into a star which preserves the property of R^P_σ of being closely related to and narrow in P.

Let $\vec{r}, \vec{s} \in \vec{S}$ be two crossing separations. An *unscrambling* of $\vec{r}, \vec{s}$ are two separations $\overrightarrow{r'}, \overrightarrow{s'} \in \vec{S}$ such that

$$\vec{r} \wedge \overleftarrow{s} \leqslant \overrightarrow{r'} \leqslant \vec{r} \text{ and } \vec{s} \wedge \overleftarrow{r} \leqslant \overrightarrow{s'} \leqslant \vec{s} \text{ and}$$
$$\text{at least one of } \overrightarrow{s'} = \vec{s} \wedge \overleftarrow{r'} \text{ and } \overrightarrow{r'} = \vec{r} \wedge \overleftarrow{s'} \text{ holds}$$

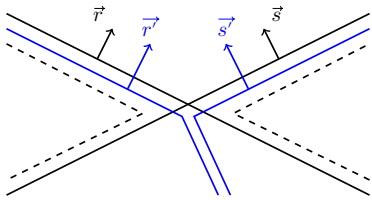

Figure 5.1. An unscrambling $\overrightarrow{r'}, \overrightarrow{s'}$ of two crossing separations $\vec{r}$ and $\vec{s}$.

(1) This generalizes a method described in [16, Chapter 4.1].

(see Figure 5.1). Note that we do not require $\overrightarrow{r'} = \vec{r} \wedge \overleftarrow{s'}$ *and* $\overrightarrow{s'} = \vec{s} \wedge \overleftarrow{r'}$ because in our applications we want to choose a suitable separation $\overrightarrow{r'}$ and then define $\overrightarrow{s'} := \vec{s} \wedge \overleftarrow{r'}$. But then, in general, $\vec{r} \wedge \overleftarrow{s'}$ need not be equal to $\overrightarrow{r'}$. However, if U is distributive, then both equations hold.

PROPOSITION 5.4. — *Let $\vec{r}, \vec{s}$ be two crossing separations in a distributive universe $\vec{U}$, and let $\overrightarrow{r'}, \overrightarrow{s'}$ be any unscrambling of them. Then $\overrightarrow{r'} = \vec{r} \wedge \overleftarrow{s'}$ and $\overrightarrow{s'} = \vec{s} \wedge \overleftarrow{r'}$.*

Proof. — Since $\overrightarrow{r'}, \overrightarrow{s'}$ is an unscrambling of $\vec{r}, \vec{s}$, at least one equation has to hold by definition; by symmetry, we may assume $\overrightarrow{s'} = \vec{s} \wedge \overleftarrow{r'}$. It follows that

$$\vec{r} \wedge \overleftarrow{s'} = \vec{r} \wedge (\vec{s} \wedge \overleftarrow{r'})^* = \vec{r} \wedge (\overleftarrow{s} \vee \overrightarrow{r'}) = (\vec{r} \wedge \overleftarrow{s}) \vee (\vec{r} \wedge \overrightarrow{r'}) = (\vec{r} \wedge \overleftarrow{s}) \vee \overrightarrow{r'} = \overrightarrow{r'}$$

since $\vec{r} \wedge \overleftarrow{s} \leqslant \overrightarrow{r'} \leqslant \vec{r}$ by definition. □

If R is a set of oriented separations, we call R' an *unscrambling of* R if R' can be obtained from R by successively unscrambling the crossing separations in R. We remark that the unscrambling of two separations is in general not unique. Moreover, the unscrambling of a set of more than two separations might further depend on the order in which the separations were unscrambled.

We now begin with the proof of Lemma 5.3. First, we show that for every pair of crossing separations $\vec{r}, \vec{s} \in P$ which are closely related to P there exists an unscrambling of them which is closely related to P. In addition, we prove that if r, s are nested with σ, then this unscrambling can be chosen to be nested with σ. As the set R^P_σ of all the maximal separations in P_σ is closely related to P by Proposition 3.1, this allows us to turn R^P_σ into a star σ' which is still closely related to P and nested with σ, by successively unscrambling the crossing separations in R^P_σ.

We then show that this star σ' is not just closely related to P and nested with σ but in fact also narrow in P. For this, we prove that R^σ_P is narrow in P, and that every unscrambling of a set of separations which is narrow in P is still narrow in P if S is contained in a distributive universe. We remark that this is in fact the only point where we need the assumption of Lemma 5.3 that the universe in which S lies is distributive.

Lastly, we show that every star in P_σ which contains the least number of elements among all stars in P_σ that are closely related to and narrow in P is near-maximal in P.

LEMMA 5.5. — *Let $\vec{S}$ be a submodular separation system, and let $\vec{r}, \vec{s} \in \vec{S}$ be two crossing separations which are closely related to some profile P*

of S. Further, let $\sigma \subseteq P$ be a star of good separations that is nested with r and s, and let $\overrightarrow{r'} \in P_\sigma$ be a separation which is minimal under all separations in P_σ that are closely related to P and greater or equal to $\vec{r} \wedge \overleftarrow{s}$. Then $\vec{s} \wedge \overleftarrow{r'}$ is an element of $\vec{S}$ and closely related to P.

By definition, $\overrightarrow{r'}, \vec{s} \wedge \overleftarrow{r'}$ is an unscrambling of $\vec{r}, \vec{s}$, and it is nested with σ by Lemma 2.1. Since $\vec{r}$ is a candidate for $\overrightarrow{r'}$, Lemma 5.5 implies that there exists an unscrambling of $\vec{r}, \vec{s}$ which is closely related to P and nested with σ.

Note that for separation systems of the form S_k inside some submodular universe, Erde showed a related statement which, similar as to Lemma 5.5, implies the existence of unscramblings that preserve the property of being closely related to some profile, but without the additional assertion of Lemma 5.5 that the unscrambling can be chosen to be nested with σ [16, Lemma 4.3]. There, the choice of the unscrambling relied on the existence of an order function.

Proof of Lemma 5.5. — Let $\overrightarrow{s'} \in P_\sigma$ be a separation which is minimal among all separations in P_σ that are closely related to P and greater or equal to $\vec{s} \wedge \overleftarrow{r'}$. For this, observe that $\vec{s}$ is a candidate for $\overrightarrow{s'}$. If $\overrightarrow{s'} = \vec{s} \wedge \overleftarrow{r'}$, we are done, so suppose that $\vec{s} \wedge \overleftarrow{r'} < \overrightarrow{s'}$. By the submodularity of $\vec{S}$, at least one of $\overrightarrow{r'} \wedge \overleftarrow{s'}$ and $\overleftarrow{r'} \wedge \overrightarrow{s'}$ is an element of $\vec{S}$; by symmetry[(2)] we may assume that $\overrightarrow{s''} := \overleftarrow{r'} \wedge \overrightarrow{s'} \in \vec{S}$. Note that $\overrightarrow{s''} = \overleftarrow{r'} \wedge \overrightarrow{s'} \leqslant \overleftarrow{r'} \wedge \vec{s} = \overleftarrow{r'} \wedge (\overleftarrow{r'} \wedge \vec{s}) \leqslant \overleftarrow{r'} \wedge \overrightarrow{s'}$, and thus $\overrightarrow{s''} = \overleftarrow{r'} \wedge \vec{s}$.

By the choice of $\overrightarrow{s'}$, we have that $\overrightarrow{s''}$ is not closely related to P, and so, by Proposition 3.2, there is a separation $\vec{t} \leqslant \overrightarrow{s'} \in P$ such that $\overrightarrow{s''} \wedge \vec{t} \notin \vec{S}$. We take a separation $\vec{t} \in \vec{S}$ which is minimal with that property. Since $\overleftarrow{r'} \wedge \vec{t} = \overleftarrow{r'} \wedge (\overrightarrow{s'} \wedge \vec{t}) = (\overleftarrow{r'} \wedge \overrightarrow{s'}) \wedge \vec{t} = \overrightarrow{s''} \wedge \vec{t} \notin \vec{S}$ by assumption, it follows, by the submodularity of $\vec{S}$, that $\overrightarrow{r''} := \overrightarrow{r'} \wedge \overleftarrow{t} \in \vec{S}$. Note that $\overleftarrow{s} \leqslant \overleftarrow{t}$ by the choice of $\vec{t}$ and $\vec{r} \wedge \overleftarrow{s} \leqslant \overrightarrow{r'}$ by assumption, and thus $\vec{r} \wedge \overleftarrow{s} \leqslant \overrightarrow{r''}$.

Since σ is good by assumption, we can fix, for every $\vec{y} \in \sigma$, some profile Q_y of S which has the property that $\overleftarrow{y}$ is closely related to it. We then set

$$\rho := \Big\{ \vec{y} \in \sigma : \overleftarrow{r''} \in Q_y \Big\}.$$

Now first assume that $\overrightarrow{r''}$ is closely related to P and nested with ρ. Then the assertion follows: By assumption, the star $\sigma' := \{\vec{y} \in \sigma : y \text{ crosses } r''\}$ is a subset of $\sigma \setminus \rho$, and hence Proposition 3.3 yields that the separation

(2) Note that the situation itself is not completely symmetric in r' and s' since we chose s' depending on r'. However, if $\overrightarrow{r'} \wedge \overleftarrow{s'} \in \vec{S}$, then swapping r' and s' in the remainder of this proof does not change anything.

$\vec{z} := \overrightarrow{r''} \wedge \bigwedge_{\overline{y} \in \sigma'} \overline{y}$ is contained in $\vec{S}$ and closely related to P. Moreover, as $\vec{r} \wedge \overleftarrow{s} \leqslant \overrightarrow{r''}$ and as $\vec{r} \wedge \overleftarrow{s}$ is nested with σ by assumption and Lemma 2.1, we have $\vec{r} \wedge \overleftarrow{s} \leqslant \vec{z}$. But since $\vec{z} \leqslant \overrightarrow{r''} < \vec{r}$, this contradicts the minimal choice of $\vec{r}$.

Therefore, it suffices to show that $\overrightarrow{r''}$ is closely related to P and nested with ρ. We first show the former. For this, suppose for a contradiction that $\overrightarrow{r''}$ is not closely related to P. By Proposition 3.2, this is witnessed by a separation $\vec{t} \leqslant \vec{r}$ in that $\overrightarrow{r''} \wedge \vec{t} \notin \vec{S}$. As $\overrightarrow{r''} \wedge \vec{t} = (\vec{r} \wedge \overleftarrow{t}) \wedge \vec{t} = \overleftarrow{t} \wedge (\vec{r} \wedge \vec{t}) = \overleftarrow{t} \wedge \vec{t}$, it follows, by the submodularity of $\vec{S}$, that $\vec{t} \wedge \overleftarrow{t} \in \vec{S}$. But this contradicts the minimal choice of $\vec{t}$: We have $\vec{t} \wedge \overleftarrow{t} < \vec{t}$ but still $\overrightarrow{s''} \wedge (\vec{t} \wedge \overleftarrow{t}) = (\overrightarrow{s''} \wedge \overleftarrow{t}) \wedge \vec{t} = \overrightarrow{s''} \wedge \vec{t} \notin \vec{S}$ as $\overrightarrow{s''} \leqslant \overleftarrow{r'} \leqslant \overleftarrow{t}$. Therefore, $\overrightarrow{r''}$ is closely related to P.

To conclude the proof, we are left to show that $\overrightarrow{r''}$ is nested with ρ. As $\overrightarrow{r''} = \vec{r'} \wedge \overleftarrow{t}$, and r' is nested with $\sigma \supseteq \rho$ by assumption, it suffices, by Lemma 2.1, to show that t is nested with ρ. For this, suppose for a contradiction that t crosses a separation $\overline{y} \in \rho$. Then $\overrightarrow{s''} \leqslant \overline{y}$. Indeed, since $\overrightarrow{s''} \leqslant \overleftarrow{r''}$ and y crosses r'' but not s'', we can only have $\overrightarrow{s''} \leqslant \overline{y}$ or $\overrightarrow{s''} \leqslant \overline{y}$; but the latter case is not possible as $\overrightarrow{s''} \not\leqslant \vec{x}$ for all $\vec{x} \in \sigma$ by construction. Further, as $\vec{t} \leqslant \overleftarrow{r''} \in Q_y$ and $\overline{y}$ is closely related to Q_y, it follows that $\vec{t} \wedge \overline{y} \in \vec{S}$. But this again contradicts the minimal choice of $\vec{t}$: We have $\vec{t} \wedge \overline{y} < \vec{t}$ and still $\overrightarrow{s''} \wedge (\vec{t} \wedge \overline{y}) = (\overrightarrow{s''} \wedge \overline{y}) \wedge \vec{t} = \overrightarrow{s''} \wedge \vec{t} \notin \vec{S}$. □

We can now conclude that for every set of separations which are closely related to some profile there is an unscrambling which is again closely related to that profile.

COROLLARY 5.6. — *Let $R \subseteq \vec{S}$ be a set which is closely related to some profile P of S. Further, let $\sigma \subseteq P$ be a star of good separations and suppose that R is nested with σ. Then there is an unscrambling R' of R that is closely related to P and nested with σ. Moreover, if R contains for some $\vec{s} \in \sigma$ a separation $\vec{r}$ with $\vec{s} \leqslant \vec{r}$, then R' does so as well.*

Proof. — Let $R = \{\vec{r}_1, \dots, \vec{r}_n\}$ be some enumeration of R. We unscramble the separations in R inductively by repeatedly applying Lemma 5.5 as follows. Set $R^0 := R$, and assume that we already constructed a set $R^{k-1} := \{\vec{r}_1^{k-1}, \dots, \vec{r}_n^{k-1}\}$ which is closely related to P and nested with σ. We then pick a pair of crossing separations r_i^{k-1}, r_j^{k-1} and apply Lemma 5.5 to them, yielding an unscrambling $\vec{r}_i^k, \vec{r}_j^k$ of $\vec{r}_i^{k-1}, \vec{r}_j^{k-1}$ which is closely related to P and nested with σ. In every step, we thus replace two crossing separations $\vec{r}_i^{k-1}, \vec{r}_j^{k-1}$ with two nested separations $\vec{r}_i^k \leqslant \vec{r}_i^{k-1}$ and $\vec{r}_j^k \leqslant \vec{r}_j^{k-1}$ such that $\vec{r}_i^k \leqslant \overleftarrow{r}_j^k$, so we do not need to consider a pair of

indices in $[n]$ twice. Hence, after at most $N := \frac{1}{2}n(n-1)$ steps, this process terminates, and we obtain a set $R' := R^N$ that is an unscrambling of R, closely related to P and nested with σ.

For the "moreover"-part, suppose that there are separations $\vec{s} \in \sigma$ and $\vec{r} \in R$ such that $\vec{s} \leqslant \vec{r}$. We show by induction on $k \in [N]$ that there is a separation $\vec{r}_i^k \in R^k$ with $\vec{s} \leqslant \vec{r}_i^k$, which clearly implies the assertion for $k := N$. The base case $k = 0$ holds by assumption, so we may assume that $k > 0$ and that there exists a separation $\vec{r}_i^{k-1} \in R^{k-1}$ with $\vec{s} \leqslant \vec{r}_i^{k-1}$. If $\vec{r}_i^{k-1} = \vec{r}_i^k$, we are done, so we may assume that we unscrambled $\vec{r}_i^{k-1}$ and some crossing separation $\vec{r}_j^{k-1}$ in the k-th step. In particular, this means that $\vec{r}_j^{k-1} \not\leqslant \vec{s}$ since otherwise $\vec{r}_j^{k-1} \leqslant \vec{s} \leqslant \vec{r}_i^{k-1}$ would contradict that r_j^{k-1} and r_i^{k-1} cross. Moreover, as P is consistent and $\vec{r}_j^{k-1}, \vec{s} \in P$, we cannot have $\overleftarrow{r}_j^{k-1} \leqslant \vec{s}$. Therefore, since r_j^{k-1} and s are nested by construction, either $\vec{s} \leqslant \overleftarrow{r}_j^{k-1}$ or $\vec{s} \leqslant \vec{r}_j^{k-1}$. In the first case we have $\vec{s} \leqslant \vec{r}_i^{k-1} \wedge \overleftarrow{r}_j^{k-1} \leqslant \vec{r}_i^k$, while in the second case we either have $\vec{s} \leqslant \vec{r}_i^k$ or $\vec{s} \leqslant \vec{r}_j^k$ by the definition of an unscrambling. □

Next, we show that the set R_σ^P of all maximal separations in P_σ is narrow in P.

Proposition 5.7. — *Let P be a profile of S, and let $\sigma \subseteq P$ be a star of good separations. Further, let $R_\sigma^P \subseteq P_\sigma$ be the set of all maximal separations in P_σ. Then R_σ^P is narrow in P.*

Proof. — Let $\vec{x} \in P$ be given; we need to show that $\overleftarrow{x} \vee \bigvee_{R_\sigma^P} \vec{r}$ is co-small. For this, let $\overrightarrow{x'} \in P$ with $\vec{x} \leqslant \overrightarrow{x'}$ be maximal in P, and fix for every separation $\vec{s} \in \sigma$ a profile Q_s of S to which $\overleftarrow{s}$ is closely related. Then $\overrightarrow{x'} \in Q_s$ for every $\vec{s} \in \sigma_{x'} := \{\vec{s} \in \sigma : \vec{s} \text{ crosses } \overrightarrow{x'}\}$ since otherwise $\vec{s} \vee \overrightarrow{x'} = (\overleftarrow{s} \wedge \overleftarrow{x'}) \in \vec{S}$, and thus $\vec{s} \vee \overrightarrow{x'} \in P$ as P is a profile, which would contradict the maximality of $\overrightarrow{x'}$ in P. It follows, by Proposition 3.3, that $\vec{y} := \overrightarrow{x'} \wedge \bigwedge_{\vec{s} \in \sigma_{x'}} \overleftarrow{s}$ is an element of $\vec{S}$. By construction, y is nested with σ, and thus P_σ contains an orientation of y; by consistency, we have $\overrightarrow{y} \in P_\sigma$. This implies that there is a separation $\overrightarrow{y'} \in R_\sigma^P$ with $\overrightarrow{y} \leqslant \overrightarrow{y'}$. Thus,

$$\overleftarrow{x} \vee \bigvee_{\vec{r} \in R_\sigma^P} \vec{r} \geqslant \overleftarrow{x'} \vee \bigvee_{\vec{r} \in R_\sigma^P} \vec{r} = \overleftarrow{x'} \vee \overrightarrow{y'} \vee \bigvee_{\vec{r} \in R_\sigma^P} \vec{r} \geqslant \overleftarrow{x'} \vee \left(\overrightarrow{x'} \wedge \bigwedge_{\vec{s} \in \sigma_{x'}} \overleftarrow{s} \right) \vee \bigvee_{\vec{r} \in R_\sigma^P} \vec{r}$$

$$\geqslant \overleftarrow{x'} \vee \left(\overrightarrow{x'} \wedge \bigwedge_{\vec{s} \in \sigma_{x'}} \overleftarrow{s} \right) \vee \bigvee_{\vec{s} \in \sigma_{x'}} \vec{s} = \left(\overrightarrow{x'} \wedge \bigwedge_{\vec{s} \in \sigma_{x'}} \overleftarrow{s} \right)^* \vee \left(\overrightarrow{x'} \wedge \bigwedge_{\vec{s} \in \sigma_{x'}} \overleftarrow{s} \right)$$

is co-small where the second last inequality holds since for every $\vec{s} \in \sigma_{x'}$ there is some $\vec{r} \in R_\sigma^P$ with $\vec{s} \leqslant \vec{r}$. □

The next lemma shows that every unscrambling of R_σ^P is narrow in P.

LEMMA 5.8. — *Let $\vec{S}$ be a submodular separation system inside some distributive universe, and let P be a profile of S. Further, let $R \subseteq P$, and let R' be any unscrambling of R. If R is narrow in P, then R' is narrow in P.*

Proof. — Since R' is an unscrambling of R, there is a sequence $R = R_0, R_1, \dots, R_n = R'$ of sets $R_i \subseteq \vec{S}$ where R_{i+1} is obtained from R_i by replacing two crossing separations in R_i with an unscrambling of them.

We show by induction on $i \in \{0, \dots n\}$ that $\overleftarrow{x} \vee \bigvee_{\vec{s} \in R_i} \vec{s}$ is co-small for every $\vec{x} \in P$. For $R_n = R'$ this yields the claim. The base case is clear since $R_0 := R$ is narrow in P by assumption, so let $i \geqslant 0$, and assume that R_i is narrow in P. Further, let $\vec{r}, \vec{s} \in R_i$ be the separations that are unscrambled in the $(i+1)$-st step. Recall that by Proposition 5.4 the unscrambling of two separations $\vec{r}, \vec{s}$ in a distributive universe yields two separations $\overrightarrow{r'} = \vec{r} \wedge \overleftarrow{s'}$ and $\overrightarrow{s'} = \vec{s} \wedge \overleftarrow{r'}$. We need to show that for every $\vec{x} \in P$ the separation

$$\begin{aligned}
&\overleftarrow{x} \vee \overrightarrow{r'} \vee (\vec{s} \wedge \overleftarrow{r'}) \vee \bigvee_{\vec{t} \in R_i \setminus \{\vec{r}, \vec{s}\}} \vec{t} \\
&= \Big(\overleftarrow{x} \vee \overrightarrow{r'} \vee \vec{s} \vee \bigvee_{\vec{t} \in R_i \setminus \{\vec{r}, \vec{s}\}} \vec{t}\Big) \wedge \Big(\overleftarrow{x} \vee \overrightarrow{r'} \vee \overleftarrow{r'} \vee \bigvee_{\vec{t} \in R_i \setminus \{\vec{r}, \vec{s}\}} \vec{t}\Big) =: \vec{y} \wedge \vec{z}
\end{aligned}$$

is co-small. For this, observe that for every pair of separations $\vec{u}, \vec{w} \in \vec{S}$ we have that $\vec{u} \wedge \vec{w}$ is co-small whenever $\vec{u}$ and $\vec{w}$ are co-small and $\vec{u} \geqslant \overleftarrow{w}$ and $\vec{w} \geqslant \overleftarrow{u}$ since then $\vec{u} \wedge \vec{w} \geqslant \overleftarrow{u}, \overleftarrow{w}$ and thus $\vec{u} \wedge \vec{w} \geqslant \overleftarrow{u} \vee \overleftarrow{w}$.

Now $\vec{z}$ is clearly co-small since $\overrightarrow{r'} \vee \overleftarrow{r'}$ is. Further, it holds that $\vec{y} \geqslant \overrightarrow{r'} \geqslant \overleftarrow{z}$ and $\vec{z} \geqslant \overleftarrow{r'} \geqslant \overleftarrow{y}$. Therefore, we are left to show that $\vec{y}$ is co-small. For this, we rearrange $\vec{y}$ to

$$\begin{aligned}
\vec{y} &= \Big(\overleftarrow{x} \vee (\vec{r} \wedge \overleftarrow{s'}) \vee \vec{s} \vee \bigvee_{\vec{t} \in R_i \setminus \{\vec{r}, \vec{s}\}} \vec{t}\Big) \\
&= \Big(\overleftarrow{x} \vee \vec{r} \vee \vec{s} \vee \bigvee_{\vec{t} \in R_i \setminus \{\vec{r}, \vec{s}\}} \vec{t}\Big) \wedge \Big(\overleftarrow{x} \vee \overleftarrow{s'} \vee \vec{s} \vee \bigvee_{\vec{t} \in R_i \setminus \{\vec{r}, \vec{s}\}} \vec{t}\Big) =: \vec{a} \wedge \vec{b},
\end{aligned}$$

where $\vec{a}$ is co-small by the induction hypothesis and $\vec{b}$ is co-small since $\overleftarrow{s'} \vee \vec{s}$ is. Further, $\vec{a} \geqslant \vec{s} \geqslant \overrightarrow{s'} \geqslant \overleftarrow{b}$ and $\vec{b} \geqslant \overleftarrow{s'} \geqslant \overleftarrow{s} \geqslant \overleftarrow{a}$, and hence $\vec{y} = \vec{a} \wedge \vec{b}$ is co-small. Therefore, R' is narrow in P. □

We are now ready to prove Lemma 5.3:

Proof of Lemma 5.3. — By Proposition 3.1 and 5.7, R_σ^P is closely related to and narrow in P. Applying Corollary 5.6 to R_σ^P thus yields an unscrambling R' of R_σ^P which is closely related to P and nested with σ. Moreover,

by Lemma 5.8, R' is narrow in P. Then $\rho := \{\vec{r} \in R' : \vec{r} \not\leqslant \vec{s} \text{ for all } \vec{s} \in R'\}$ is a star which is closely related to P, and it is clear from the definition that ρ is also narrow in P. Further, we have $\sigma \leqslant \rho$ by the "moreover"-part of Corollary 5.6. Thus, we can take a star $\sigma' \subseteq P$ such that:

(1) $\sigma \leqslant \sigma'$ and σ' is closely related to and narrow in P;
(2) the number of elements in σ' is minimal under all stars in P that satisfy (1).

We claim that σ' is near-maximal in P, which clearly implies the assertion. For this, suppose for a contradiction that σ' is not near-maximal in P, i.e. there is a separation $\vec{z} \in P$ which is greater than at least two separations $\vec{r}, \overrightarrow{r'} \in \sigma'$. We first show that we may assume $\vec{z} \in P_\sigma$. Indeed, let $\vec{y}$ be any maximal separation in P with $\vec{z} \leqslant \vec{y}$. By the assumption on σ, we may fix for every separation $\vec{s} \in \sigma$ a profile Q_s of S such that $\overleftarrow{s}$ is closely related to Q_s. If $\overleftarrow{y} \in Q_s$ for a some $\vec{s} \in \sigma_y := \{\vec{s} \in \sigma : s \text{ crosses } y\}$, then $\vec{y} \vee \vec{s} = (\overleftarrow{y} \wedge \overleftarrow{s})^*$ is an element of $\vec{S}$ and thus contained in P because P is a profile, which contradicts the maximality of $\vec{y}$ in P. Thus, $\overleftarrow{y} \in Q_s$ for all $\vec{s} \in \sigma_y$. Then Proposition 3.3 implies that $\overrightarrow{y'} := \vec{y} \wedge \bigwedge_{\vec{s} \in \sigma_y} \overleftarrow{s}$ is an element of $\vec{S}$. By Lemma 2.1, y' is nested with σ, and thus $\overrightarrow{y'} \in P_\sigma$. Moreover, since σ is a star, we still have $\vec{r}, \overrightarrow{r'} \leqslant \overrightarrow{y'}$, and thus $\overrightarrow{y'}$ is greater than at least two separations in σ.

Therefore, there is a separation in P_σ which is greater than at least two separations in σ, and we let $\vec{x} \in P_\sigma$ be any maximal such separation. Set $R := \{\vec{s} \in \sigma' : \vec{s} \not\leqslant \vec{x}\} \cup \{\vec{x}\}$. By construction, we have $\bigvee_{\sigma'} \vec{s} \leqslant \bigvee_R \vec{s}$, which implies that R is narrow in P. Moreover, $\vec{x}$ is closely related to P by Proposition 3.1. Applying Corollary 5.6 to R yields a star σ'' which is closely related to P and narrow in P by Lemma 5.8. Moreover, $\sigma \leqslant \sigma''$ by the "moreover"-part of Corollary 5.6, and thus σ'' satisfies (1). But since $\vec{x}$ was greater than at least two separations in σ', it follows that $|\sigma''| = |R| < |\sigma'|$, which contradicts that σ' satisfies (2). Hence, σ' is near-maximal in P. □

As mentioned earlier, we only needed the assumption that S lies inside a *distributive* universe in order to prove that the unscrambling of a set which is narrow in some profile is still narrow in that profile. Therefore, along the lines of the proof of Theorem 1.3, one can show that even if S lies inside some non-distributive universe, it is still possible to refine any nested set $\widetilde{N} \subseteq S$ that satisfies the premise of Theorem 1.3. The refinement N of $\widetilde{N}$ then has the property that all its inessential nodes are stars in $\mathcal{F}$, and all its essential nodes are unscramblings of the maximal separations in the tangle living at that node.

As the essential nodes of N are such unscramblings, they should still be intuitively "close" to the tangle they are home to. However, since the assumption that the universe is distributive was crucial for the proof of Lemma 5.8, we cannot say much about the "closeness" in terms of the partial order on $\vec{S}$. It thus remains an open problem whether Lemma 5.2, and hence Theorem 1.3, still hold in separation systems inside universes that are not distributive.

BIBLIOGRAPHY

[1] S. Albrechtsen, Refining trees of tangles in abstract separation systems: inessential parts, extended version, arXiv only extended version of [2], 2023.

[2] ———, Refining trees of tangles in abstract separation systems: inessential parts, *Comb. Theory* **4** (2024), no. 1, article no. 17 (27 pages), https://doi.org/10.5070/C64163854.

[3] ———, Refining tree-decompositions so that they display the k-blocks, *J. Graph Theory* **109** (2025), no. 3, p. 310–314, https://doi.org/10.1002/jgt.23230.

[4] J. Carmesin, R. Diestel, M. Hamann & F. Hundertmark, Canonical tree-decompositions of finite graphs I. Existence and algorithms, *J. Comb. Theory, Ser. B* **116** (2016), p. 1–24, https://doi.org/10.1016/j.jctb.2014.04.001.

[5] ———, Canonical tree-decompositions of finite graphs II. Essential parts, *J. Comb. Theory, Ser. B* **118** (2016), p. 268–283, https://doi.org/10.1016/j.jctb.2014.12.009.

[6] J. Carmesin & J. P. Gollin, Canonical tree-decompositions of a graph that display its k-blocks, *J. Comb. Theory, Ser. B* **122** (2017), p. 1–20, https://doi.org/10.1016/j.jctb.2016.05.001.

[7] R. Diestel, *Graph Theory*, 5th ed., Springer, 2017.

[8] ———, Abstract separation systems, *Order* **35** (2018), p. 157–170, https://doi.org/10.1007/s11083-017-9424-5.

[9] ———, Tree sets, *Order* **35** (2018), p. 171–192, https://doi.org/10.1007/s11083-017-9425-4.

[10] R. Diestel, P. Eberenz & J. Erde, Duality theorems for blocks and tangles in graphs, *SIAM J. Discrete Math.* **31** (2017), no. 3, p. 1514–1528, https://doi.org/10.1137/16M1077763.

[11] R. Diestel, C. Elbracht & R. W. Jacobs, Point sets and functions inducing tangles of set separations, *J. Comb.* **15** (2024), no. 3, p. 283–306, https://doi.org/10.4310/JOC.240907010505.

[12] R. Diestel, J. Erde & D. Weissauer, Structural submodularity and tangles in abstract separation systems, *J. Comb. Theory, Ser. A* **167C** (2019), p. 155–180, https://doi.org/10.1016/j.jcta.2019.05.001.

[13] R. Diestel, F. Hundertmark & S. Lemanczyk, Profiles of separations: in graphs, matroids, and beyond, *Combinatorica* **39** (2019), no. 1, p. 37–75, https://doi.org/10.1007/s00493-017-3595-y.

[14] R. Diestel & S.-I. Oum, Tangle-tree duality in graphs, matroids and beyond, *Combinatorica* **39** (2019), p. 879–910, https://doi.org/10.1007/s00493-019-3798-5.

[15] ———, Tangle-tree duality in abstract separation systems, *Adv. Math.* **377** (2021), article no. 107470 (25 pages), https://doi.org/10.1016/j.aim.2020.107470.

[16] J. Erde, Refining a tree-decomposition which distinguishes tangles, *SIAM J. Discrete Math.* **31** (2017), p. 1529–1551, https://doi.org/10.1137/16M1059539.

[17] N. Robertson & P. D. Seymour, Graph Minors. X. Obstructions to Tree-Decomposition, *J. Comb. Theory, Ser. B* **52** (1991), p. 153–190, https://doi.org/10.1016/0095-8956(91)90061-N.

Manuscript received 25th March 2024,
accepted 10th March 2026.

Sandra ALBRECHTSEN
University of Hamburg
Department of Mathematics
Bundestrasse 55, 20146 Hamburg (Germany)
sandra@albrechtsen-mail.de

© The author, 2026

This article is licensed under the Creative Commons Attribution 4.0 International License.
http://creativecommons.org/licenses/by/4.0/

Innovations in Graph Theory
is a member of the Mersenne
Center for Open Publishing
ISSN: 3050-743X