

Trees of tangles in abstract separation systems

Christian Elbracht Jakob Kneip Maximilian Teegen

20th September 2019

We prove canonical and non-canonical tree-of-tangles theorems for abstract separation systems that are merely structurally submodular. Our results imply all known tree-of-tangles theorems for graphs, matroids and abstract separation systems with submodular order functions, with greatly simplified and shortened proofs.

1 Introduction

Tangles were introduced in [20] by Robertson and Seymour as part of their graph minor project. Since then tangles have become one of the central objects of study in the field of graph minor theory. The general idea of tangles is to capture highly connected substructures in graphs not by listing the vertices and edges belonging to that structure but instead, indirectly, by listing for every low-order separation of the graph on which side of that separation the highly connected structure lies. In this way, the tangle ‘points towards’ the dense structure rather than describing it explicitly. The paradigm shift brought about by Robertson and Seymour, then, was to view tangles themselves as highly cohesive objects, and working directly with these sets of oriented separations rather than with any substructures possibly captured by them.

The central theorem in the theory of tangles, which was established by Robertson and Seymour together with the notion of tangles, is the following:

Theorem 1.1 ([20]). *Every graph has a tree-decomposition displaying its maximal tangles.*

Theorem 1.1 roughly says that the highly cohesive regions in a graph are arranged in a tree-like structure. The ‘maximal’ in Theorem 1.1 relates to the order of the tangles: the tree-decomposition found by Theorem 1.1 displays the graph’s tangles at every level of coarseness.

The original proof of Theorem 1.1 by Robertson and Seymour in [20] is fairly involved and uses as tools multiple non-trivial results about separations in graphs, for instance, the existence of certain ‘tie-breaker’ functions. Since then, the theory of tangles has moved on considerably, and shorter and more elementary proofs of Theorem 1.1 have

been found. The shortest proof to date is due to Carmesin [3, 6], who utilises the fact that the separations needed for the tree-decomposition in Theorem 1.1 behave well under appropriately defined joins and meets when taken to be of minimal order.

Carmesin, Diestel, Hundertmark, and Stein established the following strengthening of Theorem 1.1:

Theorem 1.2 ([4]). *Every graph has a canonical tree-decomposition displaying its maximal tangles.*

Here, ‘canonical’ means that every automorphism of the graph acts on the decomposition tree. In other words, Theorem 1.2 uses only invariants of the graph—in particular no tie-breaker—to find the desired tree-decomposition.

A careful analysis of [4]’s proof of Theorem 1.2 conducted in [9] resulted in another shift of paradigm, similarly to the shift brought about by [20]. Much in the same way as tangles made it possible indirectly to capture substructures in graphs that were traditionally described more directly by sets of vertices or edges, and to treat them in a unified framework, it turned out that tangles themselves could be described, unlike in their definition given by Robertson and Seymour in [20], without reference to vertices or edges.

Indeed, the only information needed about a graph’s tangles to prove Theorem 1.2 is how its separations relate to each other, that is, which separations are nested or cross. Formally, the separations of the graph are turned into a poset together with an involution, and all subsequent tools and theorems in [9] are then formulated for these posets. This new notion of ‘abstract tangles’ yielded not only a cleaner proof of Theorem 1.2 in [9], but also made the theory of tangles applicable to a wider range of combinatorial structures.

This novel way of working with so-called ‘abstract separation systems’ is summarised in [7], and has yielded multiple generalisations and strengthenings of various theorems in the theory of tangles in both the finite (see [2, 8–11, 13]) and the infinite (see [8, 12, 16, 17]) setting. In this generalised framework tree-decompositions and tangles of graphs are generalised to nested sets of separations and ‘profiles’, respectively. Working with abstract separation systems rather than with graphs makes many of the results in this new theory of tangles applicable to give notions of highly cohesive substructures in settings other than just graphs, such as in matroids or in image segmentation ([14, 15]).

However one condition not expressed in terms of relations between the separations remained in use throughout the series of abstractions of Theorem 1.1 implemented in [4] and [9]: all of these works assumed that the separation systems of interest came with a submodular order function. Likewise, Carmesin’s short proof of Theorem 1.1 in [3] also leverages the fact that the order of separations of graphs is a submodular function.

This last non-structural aspect of tree-of-tangles theorems was disposed of in [11]: in that paper Diestel, Erde, and Weißbauer replaced the order function with a purely structural notion of submodularity which can be expressed solely in terms of the lattice structure of the separation system. In doing so they established the most general and widely applicable variant of Theorem 1.1 to date:

Theorem 1.3 ([11, Theorem 6]). *Let \vec{S} be a structurally submodular separation system and \mathcal{P} a set of profiles of S . Then S contains a nested set that distinguishes \mathcal{P} .*

The relevant separation systems in graphs are all ‘structurally submodular’, and therefore Theorem 1.3 still applies to tangles in graphs. On the other hand there are separation systems that are structurally submodular but cannot be represented by graph separations ([2]). In particular Theorem 1.3 can also be applied to separation systems which, unlike separations in graphs, do not come with any order function, such as arbitrary bipartitions of sets. This is a marked step forward from its predecessor Theorem 1.1, whose original proof made heavy use of the order of particular separations.

However there is a trade-off involved in Theorem 1.3’s wider applicability: it does not imply Theorem 1.1. Indeed, Theorem 1.3 applied to a graph produces a tree-decomposition which displays just the graph’s k -tangles for arbitrary but fixed k . This is a significant weakening of Theorem 1.1, which finds a decomposition displaying the graph’s maximal tangles for all tangles orders simultaneously. Moreover, the tree-decomposition found by Theorem 1.1 is *efficient* in the sense that for every pair of tangles distinguished by the tree-decomposition, the separation in the decomposition distinguishing that pair of tangles is of the lowest possible order. Since Theorem 1.3 makes only structural assumptions so as to be applicable to separation systems without any order function, Theorem 1.3 cannot guarantee that the separations used by the nested set to distinguish a particular pair of tangles are of minimal order.

In this paper we bridge the gap between Theorem 1.1 and Theorem 1.3 by establishing the following tree-of-tangles theorem which combines the upsides of both Theorem 1.1 and Theorem 1.3, i.e., which is as widely applicable as Theorem 1.3 while still being as powerful and efficient as Theorem 1.1 when applied to tangles in graphs:

Theorem 1.4. *If $\mathcal{S} = (S_1, \dots, S_n)$ is a compatible sequence of structurally submodular separation systems inside a universe U , and \mathcal{P} is a robust set of profiles in \mathcal{S} , then there is a nested set N of separations in U which efficiently distinguishes all the distinguishable profiles in \mathcal{P} .*

Theorem 1.4 includes Theorem 1.3 by taking a sequence of just one separation system, and it implies Theorem 1.1 by taking as separation systems S_k the sets of all separations of order $< k$ of the given graph; the resulting nested set is the set of separations of the desired tree-decomposition.

The nested set N found by Theorem 1.4 has to contain for every pair of profiles in \mathcal{P} a separation from that pair’s ‘candidate set’ of all those separations which (efficiently) distinguish that pair of profiles. Thus, to prove Theorem 1.4, it suffices to show that one can pick an element from each of these ‘candidate sets’ in a nested way.

As it turns out, there is a very simple and purely structural requirement of the way these ‘candidate sets’ interact with each other which guarantees that it is possible to pick such a nested set:

Theorem 1.5 (Splinter theorem). *Let U be a universe of separations and $\mathcal{A} = (A_i)_{i \leq n}$ a family of subsets of U . If \mathcal{A} splinters then we can pick an element a_i from each A_i so that $\{a_1, \dots, a_n\}$ is nested.*

Theorem 1.5, in a sense, represents yet another step of abstraction in the theory

of tangles: rather than working with the profiles themselves it works with the sets of separations distinguishing a given pair of profiles.

Theorem 1.5 not only implies Theorem 1.4, but can also be used to prove Theorem 1.1 and Theorem 1.3 directly. In fact Theorem 1.5 has a remarkably short proof (as we shall see in Section 3), making it the shortest available proof of Theorem 1.1 so far (see Section 4.1). Moreover, the premise in Theorem 1.5 is straightforward to check, and Theorem 1.5 itself does not make reference to tangles or any specific implementations of them. As a result Theorem 1.5 can be used in many different settings, implying variations of Theorem 1.1 in a multitude of contexts. For example, after deriving in Section 4 Theorem 1.1, Theorem 1.3, and Theorem 1.4 from Theorem 1.5, we use Theorem 1.5 to establish a new tree-of-tangles theorem in the setting of clique separations.

Since Theorem 1.5 does not yield a canonical set of separations we cannot deduce Theorem 1.2 from it. We fix this in Section 5 by establishing a version of Theorem 1.5 which does give a canonical nested set, albeit under slightly stronger assumptions:

Theorem 1.6 (Canonical splinter theorem). *Let U be a universe of separations and let $\mathcal{A} = (A_i \mid i \in I)$ be a collection of subsets of U that splinters hierarchically with respect to a partial order \preceq on I . Then there exists a nested set $N = N(\mathcal{A})$ meeting every A_i in \mathcal{A} .*

Moreover, $N(\mathcal{A})$ is canonical: if φ is an isomorphism of separation systems between $\bigcup_{i \in I} \vec{A}_i$ and a subset of some universe U' such that the family $\varphi(\mathcal{A}) := (\varphi(A_i) \mid i \in I)$ splinters hierarchically with respect to \preceq , then $N(\varphi(\mathcal{A})) = \varphi(N(\mathcal{A}))$.

We make use of Theorem 1.6 in Section 6 to obtain a new shortest proof of Theorem 1.2 and to extend Theorem 1.2 to two natural types of separations whose structural submodularity does not come from a submodular order function: clique separations, and circle separations.

All graph-theoretic terms and notation used but not explained here can be found in [6].

2 Separation systems and profiles

Our terms and notation regarding separation systems follow those of [7], whereto we refer the reader for an in-depth introduction to abstract separation systems. In addition we use terms and definitions from [9], which introduced the notion of profiles. For the convenience of the reader we shall offer a brief introduction to separation systems and profiles in this section. In addition to this we will work with a slightly more general notion of k -profiles than that used in [9]; we shall indicate below where we deviate from the usual definitions.

A *separation system* is a poset \vec{S} with an order-reversing involution $*$. If an element of \vec{S} is introduced as \vec{s} , we denote its inverse under the involution by $\vec{s}^* := (\vec{s})^*$. The involution on \vec{S} being order-reversing then means that $\vec{r} \leq \vec{s}$ if and only if $\vec{r}^* \geq \vec{s}^*$ for all $\vec{r}, \vec{s} \in \vec{S}$. The elements of \vec{S} are called (oriented) *separations*. In this paper all separation systems (and graphs) are assumed to be finite.

Given a separation $\vec{s} \in \vec{S}$ we write s for the set $\{\vec{s}, \vec{s}^*\}$, and call s the underlying *unoriented separation* of \vec{s} (and of \vec{s}^*). Given an unoriented separation s , we call \vec{s} and

\bar{s} the two *orientations* of s . If \vec{S} is a separation system we write S for the set of all s with $\vec{s} \in \vec{S}$. For clarity of writing we may informally use terms that are defined only for unoriented separations also when referring to their orientations, and vice-versa. Likewise, we use the term ‘separation’ for both oriented and unoriented separations.

Two separations r and s are *nested* if they have orientations \vec{r} and \vec{s} such that $\vec{r} \leq \vec{s}$; otherwise r and s *cross*. A set of separations is nested if all its elements are pairwise nested.

A *universe of separations* is a separation system \vec{U} in which the poset comes with join and meet operators \vee and \wedge which make it a lattice. Note that in a universe DeMorgan’s law holds, i.e., $(\vec{r} \vee \vec{s})^* = \vec{r} \wedge \vec{s}$. Given two separations r and s their *corner separations* are the separations of the form $\vec{r} \vee \vec{s}$ (or their underlying unoriented separations), where \vec{r} and \vec{s} range over all possible orientations of r and s .

A separation $\vec{r} \in \vec{S}$ is *trivial* if $\vec{r} < \vec{s}$ and $\vec{r} < \vec{s}$ for some $\vec{s} \in \vec{S}$. An unoriented separation is trivial if it has a trivial orientation. If $T \subseteq S$ is a set of non-trivial separations that are pairwise nested, then T is a *tree set*.

Separations $\vec{s} \in \vec{S}$ with $\vec{s} \leq \bar{s}$ are called *small*. A set of oriented separations is *regular* if it contains no small separation, and a set of unoriented separations is *regular* if none of its elements has an orientation that is small. Note that if \vec{s} is trivial then \bar{s} is small; the converse of this is not true in general.

For a universe \vec{U} , an *order function* on \vec{U} is a function $|\cdot|: \vec{U} \rightarrow \mathbb{R}_{\geq 0}$ with $|\vec{s}| = |\bar{s}|$ for all $\vec{s} \in \vec{U}$. Given such an order function we let $|s| := |\vec{s}|$. The universe \vec{U} is called *submodular* if the order function is submodular, i.e., if

$$|\vec{r}| + |\vec{s}| \geq |\vec{r} \vee \vec{s}| + |\vec{r} \wedge \vec{s}| \quad \forall \vec{r}, \vec{s} \in \vec{U}.$$

A separation system $\vec{S} \subseteq \vec{U}$ inside a universe is called (*structurally*) *submodular* if for any $\vec{r}, \vec{s} \in \vec{S}$ at least one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ also lies in \vec{S} .

An *orientation* of a set S of unoriented separations is a set consisting of precisely one orientation, \vec{s} or \bar{s} , for every $s \in S$. An orientation O is *consistent* if there are no $\vec{r}, \vec{s} \in O$ with $\vec{r} < \vec{s}$. Two orientations O and O' (not necessarily of the same set S) are *distinguished* by some separation s if they orient s differently; if O and O' are distinguished by some s then they are *distinguishable*. In a context where an order function is given, s distinguishes O and O' *efficiently* if it is of minimal order among all separations that distinguish O and O' .

Given a universe U and $S \subseteq U$, a *profile* of S is a consistent orientation P of S with the *profile property*:

$$\forall \vec{r}, \vec{s} \in P: (\vec{r} \wedge \vec{s}) \notin P \tag{P}$$

If the universe of separations U has an order function, then for a fixed $S \subseteq U$ we let $S_k := \{s \in S \mid |s| < k\}$. A *k-profile* (in S), then, is a profile of S_k , and a *profile in S* is a k -profile for some integer k .

The above definitions of S_k and (k -)profiles in S are a minor generalisation of the definitions used in [9]. That paper exclusively studied profiles in U , that is, always took $S = U$ in these definitions. Indeed, if no S is specified we shall implicitly take $S = U$ in

these definitions; observe that in that case the separation systems S_k are submodular if the order function on U is submodular.

Traditionally, a separation of a graph $G = (V, E)$ is a pair (A, B) of subsets of V with $A \cup B = V$ such that there is no edge in G between $A \setminus B$ and $B \setminus A$. These separations of graphs are an instance of separation systems: the set $\vec{S}(G)$ of all separations of G constitutes a separation system when equipped with the involution $(A, B)^* := (B, A)$ and the partial order where $(A, B) \leq (C, D)$ if and only if $A \subseteq C$ and $B \supseteq D$. In fact, this partial order on $\vec{S}(G)$ is a lattice with $(A, B) \vee (C, D) = (A \cup C, B \cap D)$, and $\vec{S}(G)$ becomes a submodular universe by letting the order of a separation be $|(A, B)| := |A \cap B|$. The notion of profiles in separation systems then generalises and extends the notion of tangles in graphs: every k -tangle in G is in fact a k -profile in this universe $\vec{S}(G)$ of all separations of G .

3 The Splinter Theorem

In this section we establish our first main theorem, Theorem 1.5, from which we shall derive two previously known results as well as two new flavours of tree-of-tangles theorems in Section 4. A cornerstone of the proofs of both Theorem 1.5 as well as of the two known results we shall derive from it is the following lemma:

Lemma 3.1 ([9, Lemma 2.1]). *Let U be a universe and $r, s \in U$ 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 .*

Typically, the proof of a tree-of-tangles theorem proceeds by starting with some set N of separations which distinguish some (or all) of the given tangles, and then repeatedly replacing elements r of N which cross some other element s of N with an appropriate corner separation of r and s . Lemma 3.1 is then used to show that each of these replacements makes N ‘more nested’, and thus one eventually obtains a nested set N which distinguishes all the given tangles. (See for instance the proof of Theorem 4 of [11].) Usually, in order to not reduce the set of tangles distinguished by N , one has to take special care which corner separation of two crossing r and s in N one uses for replacement; this depends on the specific properties of the tangles at hand.

Our Theorem 1.5 seeks to eliminate this careful selection of corner separations for replacement: we will show that for a family $(A_i)_{i \leq n}$ of subsets of some universe U we can find a nested set N meeting all the A_i , provided that these sets A_i have one straightforward-to-check property. This theorem will imply many of the existing tree-of-tangles theorems by taking as sets A_i the sets of separations which distinguish the i -th pair of tangles, and checking that the one assumption needed for Theorem 1.5 is met. Notably, Theorem 1.5 will make no reference at all to tangles or their specific properties. The proof of Theorem 1.5 will also utilise Lemma 3.1; however, the only assumption we need about the sets A_i is that for elements a_i and a_j of A_i and A_j , respectively, one of their four corner separations lies in either A_i or A_j . This condition will be easy to verify if one wants to deduce other tree-of-tangles theorems from Theorem 1.5. In fact, the

verification of this condition, which just asks for the existence of *some* corner separation of r and s in $A_i \cup A_j$, will usually be much more straightforward than the hands-on arguments used in the original proofs of those tree-of-tangles theorems, which for their replacement arguments often need to prove the existence of a *specific* corner separation of r and s . So let us define this condition formally.

Let U be a universe and $\mathcal{A} = (A_i)_{i \leq n}$ some family of non-empty subsets of U . We say that \mathcal{A} *splinters* if, for every pair $a_i \in A_i$ and $a_j \in A_j$ of separations, either some corner separation of a_i and a_j lies in $A_i \cup A_j$, or one of a_i and a_j lies in $A_i \cap A_j$.

Observe that a family $(A_i)_{i \leq n}$ of non-empty sets splinters if and only if for all crossing $a_i \in A_i \setminus A_j$ and $a_j \in A_j \setminus A_i$ one of their four corner separations lies in $A_i \cup A_j$: for if two separations a_i and a_j are nested, these separations themselves are corner separations of the pair a_i and a_j .

With this definition and Lemma 3.1 we are already able to state and prove our first main result:

Theorem 1.5 (Splinter theorem). *Let U be a universe of separations and $\mathcal{A} = (A_i)_{i \leq n}$ a family of subsets of U . If \mathcal{A} splinters then we can pick an element a_i from each A_i so that $\{a_1, \dots, a_n\}$ is nested.*

Proof. We proceed by induction on n . The assertion clearly holds for $n = 1$. So suppose that $n > 1$ and that the above assertion holds for all smaller values of n .

Suppose first that we can find some $a_i \in A_i$ such that a_i is nested with at least one element of A_j for each $j \neq i$. Then the assertion holds: for $j \neq i$ let A'_j be the set of those elements of A_j that are nested with a_i . Then $(A'_j \mid j \neq i)$ is a family of non-empty sets which splinters by Lemma 3.1. Thus by the induction hypothesis we can pick a nested set $\{a_j \in A'_j \mid j \neq i\}$, which together with a_i is the desired nested set.

To conclude the proof it thus suffices to find an a_i as above. To this end, we apply the induction hypothesis to A_1, \dots, A_{n-1} to obtain a nested set consisting of some a_1, \dots, a_{n-1} . Fix an arbitrary $a_n \in A_n$. For all $i < n$, if a_i itself or one of its corner separations with a_n lies in A_n , this a_i is the desired separation for the above argument. Otherwise, for each $i < n$, either a_n itself or one of its corner separations with a_i lies in A_i , in which case a_n is the desired separation for the above argument. \square

We shall see in Section 4 that this innocuous-looking theorem is actually strong enough to directly imply various existing tree-of-tangles theorems, including Theorem 1.1.

4 Applications of the Splinter Theorem

4.1 A short proof of Theorem 1.1

As a first application of Theorem 1.5 let us give a short proof of Theorem 1.1:

Theorem 1.1 ([20]). *Every graph has a tree-decomposition displaying its maximal tangles.*

Let us first recall the relevant definitions of separations and tangles in graphs:

Let $G = (V, E)$ be a graph. Then the set $\vec{S} = \vec{S}(G)$ of all separations (A, B) of G is a separation system with involution $(A, B)^* = (B, A)$ and the partial order in which $(A, B) \leq (C, D)$ if and only if $A \subseteq C$ and $B \supseteq D$. The *order* of a separation (A, B) of G is $|(A, B)| := |A \cap B|$. With this order function \vec{S} becomes a submodular universe, where $(A, B) \vee (C, D) = (A \cup C, B \cap D)$. For $(A, B) \in \vec{S}$ we write $\{A, B\}$ for the underlying unoriented separation. Furthermore we write $\vec{S}_k = \vec{S}_k(G)$ for the set of all separations of G of order $< k$.

A k -tangle in G , for an integer k , is an orientation P of \vec{S}_k with the *tangle property*:

$$\forall (A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau: G[A_1] \cup G[A_2] \cup G[A_3] \neq G \quad (\text{T})$$

A k -tangle of G is a *maximal* tangle of G if it is not the subset of some l -tangle of G for some $l > k$.

For a tree-decomposition (T, \mathcal{V}) of G and an edge xy of T let T_x and T_y denote the components of $T - xy$ containing x and y , respectively. Let U_x be the union of all bags V_z with $z \in T_x$, and similarly let U_y be the union of all bags V_z with $z \in T_y$. Then we call (U_x, U_y) the *separation induced by xy* . The *set of separations induced by (T, \mathcal{V})* is the set of all separations induced by the edges of T .

We say that a tree-decomposition (T, \mathcal{V}) of G *displays its maximal tangles* if the set of separations induced by (T, \mathcal{V}) efficiently distinguishes the set of all maximal tangles of G .

If N is a nested set of separations of G it is straightforward to find a tree-decomposition of G whose set of induced separations is precisely N (see [8, 20]). Therefore, in order to prove Theorem 1.1, it suffices to find a nested set N of separations of G which efficiently distinguishes all maximal tangles of G .

For every pair P, P' of distinct maximal tangles of G let

$$A_{P, P'} := \{ \{A, B\} \in S(G) \mid \{A, B\} \text{ efficiently distinguishes } P \text{ and } P' \}.$$

Since P and P' are not subsets of each other $A_{P, P'}$ is a non-empty set.

Let \mathcal{A} be the family of all these sets $A_{P, P'}$. A nested set of separations of G distinguishes all maximal tangles of G efficiently if and only if it contains an element of each $A_{P, P'}$. Therefore the existence of such a set, and hence Theorem 1.1, now follows directly from Theorem 1.5 once we show that \mathcal{A} splinters:

Lemma 4.1. *The family \mathcal{A} of all $A_{P, P'}$ splinters.*

Proof. Let $P \neq P'$ and $Q \neq Q'$ be two pairs of distinct maximal tangles of G and let $\{A, B\} \in A_{P, P'}$ and $\{C, D\} \in A_{Q, Q'}$ be two crossing separations. We need to show that we have either $\{A, B\} \in A_{Q, Q'}$ or $\{C, D\} \in A_{P, P'}$, or that some corner separation of $\{A, B\}$ and $\{C, D\}$ lies in $A_{P, P'} \cup A_{Q, Q'}$. By switching their roles if necessary we may assume that $|(A, B)| \leq |(C, D)|$.

Since Q and Q' both orient (C, D) , and $|(A, B)| \leq |(C, D)|$, both tangles also orient $\{A, B\}$. If Q and Q' orient $\{A, B\}$ differently, then $\{A, B\}$ distinguishes them efficiently

and hence lies in $A_{Q,Q'}$. So suppose that Q and Q' contain the same orientation of $\{A, B\}$, say, (A, B) .

By renaming them if necessary we may assume that $(C, D) \in Q$ and $(D, C) \in Q'$.

Consider the corner separation $(A \cup C, B \cap D)$. Suppose first that $|(A \cup C, B \cap D)| \leq |(C, D)|$. Then, by $(A, B), (C, D) \in Q$ and the tangle property **(T)**, Q must contain $(A \cup C, B \cap D)$. On the other hand Q' must contain its inverse $(B \cap D, A \cup C)$ since $(D, C) \in Q'$. But then this corner separation efficiently distinguishes Q and Q' and hence lies in $A_{Q,Q'}$.

Thus we may suppose that $|(A \cup C, B \cap D)| \geq |(C, D)|$. By a similar argument we may further suppose that $|(A \cup D, B \cap C)| \geq |(C, D)|$. Submodularity then yields $|(A \cap C, B \cup D)|, |(A \cap D, B \cup C)| \leq |(A, B)|$.

By switching the roles of P and P' if necessary we may assume that $(A, B) \in P$ and $(B, A) \in P'$. Then by the above inequality P must contain both $(A \cap C, B \cup D)$ and $(A \cap D, B \cup C)$, since it cannot contain either of their inverses due to $(A, B) \in P$ and the tangle property **(T)**. However, due to $(B, A) \in P'$ and the tangle property **(T)**, P' cannot contain both of $(A \cap C, B \cup D)$ and $(A \cap D, B \cup C)$. It must therefore contain the inverse of at least one of these corner separations, which then efficiently distinguishes P and P' and hence lies in $A_{P,P'}$. \square

4.2 Abstract tangles in structurally submodular separation systems

The most general, or most widely applicable, tree-of-tangles theorem published so far, in the sense of having the weakest premise, is the following:

Theorem 1.3 ([11, Theorem 6]). *Let \vec{S} be a structurally submodular separation system and \mathcal{P} a set of profiles of S . Then S contains a nested set that distinguishes \mathcal{P} .*

The price to pay in Theorem 1.3 for having the mildest set of requirements is that its assertion is also among the weakest of all tree-of-tangles theorems. For graphs, Theorem 1.3 implies only that *for any fixed k* every graph has a tree decomposition displaying its k -tangles. This is a much weaker statement than Theorem 1.1, which finds a tree-decomposition displaying the maximal k -tangles of that graph for all values of k simultaneously.

Let us show how to derive Theorem 1.3 from Theorem 1.5. For this, let \mathcal{P} be a set of profiles of a submodular separation system S , and for distinct P and P' in \mathcal{P} let

$$A_{P,P'} := \{s \in S \mid s \text{ distinguishes } P \text{ and } P'\}.$$

For proving Theorem 1.3 it suffices to show that the family $A_{\mathcal{P}} = (A_{P,P'} \mid P \neq P' \in \mathcal{P})$ splinters:

Lemma 4.2. *Given a set \mathcal{P} of profiles of a submodular separation system \vec{S} , the family $A_{\mathcal{P}} = (A_{P,P'} \mid P \neq P' \in \mathcal{P})$ splinters.*

Proof. Let $P \neq P'$ and $Q \neq Q'$ be two pairs of profiles in \mathcal{P} and let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$ be two distinct separations. We need to show that we have either $r \in A_{Q,Q'}$

or $s \in A_{P,P'}$, or that some corner separation of r and s lies in $A_{P,P'} \cup A_{Q,Q'}$. If r and s are nested then they themselves are corner separations of r and s and there is nothing to show, so let us suppose that r and s cross.

Both r and s are oriented by all four profiles $P, P', Q,$ and Q' . If r distinguishes Q and Q' , or if s distinguishes P and P' , we are done; so suppose that there are orientations \vec{r} and \vec{s} of r and s with $\vec{r} \in Q \cap Q'$ and $\vec{s} \in P \cap P'$. By possibly switching the roles of P and P' , or of Q and Q' , we may further assume that $\vec{r} \in P$ and $\vec{s} \in P'$ as well as $\vec{r} \in Q$ and $\vec{s} \in Q'$.

The submodularity of S implies that at least one of the two corner separations $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ lies in \vec{S} . We will only treat the case that $(\vec{r} \vee \vec{s}) \in \vec{S}$; the other case is symmetrical.

From the assumption that r and s cross it follows that $\vec{r} \vee \vec{s}$ is distinct from r and s as an unoriented separation. Therefore, by $\vec{r} \in P'$ and consistency, P' cannot contain $\vec{r} \vee \vec{s}$ and hence has to contain its inverse $\vec{r} \wedge \vec{s}$. On the other hand, by $\vec{r}, \vec{s} \in P$ and the profile property (P), P cannot contain the inverse of $\vec{r} \vee \vec{s}$ and thus must contain $\vec{r} \vee \vec{s}$. Now $\vec{r} \vee \vec{s}$ distinguishes P and P' and is therefore the desired corner separation in $A_{P,P'}$. \square

Let us now deduce Theorem 1.3 from Theorem 1.5.

Proof of Theorem 1.3. Let \mathcal{P} be a set of profiles of S . By Lemma 4.2 the collection $(A_{P,P'} \mid P \neq P' \in \mathcal{P})$ of subsets of S splinters. Each of the $A_{P,P'}$ is non-empty as P and P' are distinct profiles of S . Thus, by Theorem 1.5, we can pick one element from each $A_{P,P'}$ so that the set N of all these elements is a nested set of separations. It is then clear that N distinguishes all the profiles in \mathcal{P} . \square

The above way of using Theorem 1.5 to prove a tree-of-tangles theorem is archetypical, and we will use the strategy from this section as a blueprint for the applications of Theorem 1.5 in the following sections.

4.3 Profiles of separations

Theorem 1.3 from the previous section implied that every graph has, for any fixed integer k , a tree-decomposition which displays its k -tangles. However, Robertson's and Seymour's Theorem 1.1 shows that every graph has a tree-decomposition which displays all its *maximal* tangles, i.e., which distinguishes all its distinguishable tangles for all values of k simultaneously, not just for some fixed value of k . Therefore Theorem 1.3 does not imply Theorem 1.1.

Moreover, since Theorem 1.3 does not assume that the universe U it is applied to comes with an order function, Theorem 1.3 cannot say anything about the order of the separations used in the nested set to distinguish all the profiles. If the universe U , as for instance in a graph, *does* come with a submodular order function, one might ask for a nested set which not only distinguishes all the profiles given, but one which does so *efficiently*, i.e., which contains for every pair P, P' of profiles a separation of minimal order among all the separations in U which distinguish P and P' .

The following theorem satisfies both of the requirements above, and is the strongest tree-of-tangles theorem known so far:

Theorem 4.3 (Canonical tangle-tree theorem for separation universes [9, Theorem 3.6]). *Let $U = (\vec{U}, \leq, *, \vee, \wedge, | |)$ be a submodular universe of separations. Then for every robust set \mathcal{P} of profiles in U there is a nested set $T = T(\mathcal{P}) \subseteq U$ of separations such that:*

- (i) every two profiles in \mathcal{P} are efficiently distinguished by some separation in T ;
- (ii) every separation in T efficiently distinguishes a pair of profiles in \mathcal{P} ;
- (iii) for every automorphism α of \vec{U} we have $T(\mathcal{P}^\alpha) = T(\mathcal{P})^\alpha$; (canonicity)
- (iv) if all the profiles in \mathcal{P} are regular, then T is a regular tree set.

Since the definition of robustness of (a set of) profiles is rather involved we do not repeat it here. In the following proofs robustness will be used only in one place; therefore we shall use it there as a black box and refer the reader to [9] for the full definition.

Since every k -tangle of a graph is robust ([9]), Theorem 4.3 indeed implies Theorem 1.1 of Robertson and Seymour that every graph has a tree-decomposition displaying its maximal tangles (see [9, Section 4.1] for more on building tree-decompositions from nested sets of separations, and how Theorem 4.3 implies Theorem 1.1). Moreover, Theorem 4.3 improves upon Theorem 1.1 by finding a *canonical* such tree-decomposition, i.e., one which is preserved by automorphisms of the graph. Since Theorem 1.5 does not guarantee any kind of canonicity, we are not able to deduce the full Theorem 4.3 from Theorem 1.5; however, using Theorem 1.5 we will be able to find a nested set $T \subseteq U$ with the properties (i), (ii) and (iv). We shall refer to this as the *non-canonical Theorem 4.3*. (In Section 5 we shall prove a version of Theorem 1.5 which implies Theorem 4.3 in full.)

Our strategy will largely be the same as in Section 4.2. For a robust set \mathcal{P} of profiles in a submodular universe U we define for every pair P, P' of distinct profiles in \mathcal{P} the set

$$A_{P,P'} := \{a \in U \mid a \text{ distinguishes } P \text{ and } P' \text{ efficiently}\}.$$

Let $\mathcal{A}_{\mathcal{P}}$ be the family $(A_{P,P'} \mid P \neq P' \in \mathcal{P})$. The only lemma we need in order to apply Theorem 1.5 is the following:

Lemma 4.4. *For a robust set \mathcal{P} of profiles in U the family $\mathcal{A}_{\mathcal{P}}$ of the sets $A_{P,P'}$ splinters.*

Proof. Let P, P' and Q, Q' be two pairs of distinguishable profiles in \mathcal{P} and let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$ be two crossing separations. We need to show that we have either $r \in A_{Q,Q'}$ or $s \in A_{P,P'}$, or that some corner separation of r and s lies in $A_{P,P'} \cup A_{Q,Q'}$. By switching their roles if necessary we may assume that $|r| \leq |s|$.

Since Q orients all separations in U of order at most the order of s , Q contains some orientation \vec{r} of r . Similarly Q' contains some orientation of r : if $\vec{r} \in Q'$ then r distinguishes Q and Q' , and by $|r| \leq |s|$ it does so efficiently, giving $r \in A_{Q,Q'}$. So suppose that $\vec{r} \in Q'$.

If either one of the two corner separations $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \vec{s}$ has order at most the order of s , then that corner separation would distinguish Q and Q' by the profile property. In

particular, that corner separation would do so efficiently and hence lie in $A_{Q,Q'}$. Thus we may assume that both of these corner separations have order strictly larger than the order of s .

The submodularity of U now implies that both of the other two corner separations, that is, $\vec{r} \wedge \vec{s}$ and $\vec{r} \wedge \vec{s}$, have order strictly less than the order of r . Therefore both P and P' orient both of these corner separations. By possibly switching the roles of P and P' we may assume that $\vec{r} \in P$ and $\vec{r} \in P'$. Then P' contains both $\vec{r} \wedge \vec{s}$ and $\vec{r} \wedge \vec{s}$ due to consistency, since both of these corner separations are distinct from r as unoriented separations by the assumption that r and s cross.

But now the assumption that r distinguishes P and P' efficiently implies that neither of the two corner separations $\vec{r} \wedge \vec{s}$ and $\vec{r} \wedge \vec{s}$ can distinguish P and P' , since the corner separations have strictly lower order than r . Therefore P contains $\vec{r} \wedge \vec{s}$ and $\vec{r} \wedge \vec{s}$ as well. However, by $\vec{r} \in P$, this contradicts the robustness of P , which forbids exactly this configuration. \square

Let us now deduce the non-canonical Theorem 4.3 from Theorem 1.5:

Proof of the non-canonical Theorem 4.3. By Lemma 4.4 the collection $\mathcal{A}_{\mathcal{P}}$ of the sets $A_{P,P'}$ splinters. Thus by Theorem 1.5 we can pick an element from each set $A_{P,P'}$ in $\mathcal{A}_{\mathcal{P}}$ in such a way that the set T of these elements is nested. Let us show that this set T is as claimed.

For (i), let P and P' be two profiles in \mathcal{P} . As T meets the set $A_{P,P'}$, some element of T distinguishes P and P' by definition of $A_{P,P'}$.

For (ii), observe that every element of T lies in some $A_{P,P'}$ and hence distinguishes a pair of profiles in \mathcal{P} efficiently.

Finally, (iv) follows from the fact that all sets $A_{P,P'}$ in $\mathcal{A}_{\mathcal{P}}$ are regular if every profiles in \mathcal{P} is regular, which implies that T is a regular tree set in that case. \square

4.4 Sequences of submodular separation systems

Let us, once more, compare Theorem 1.3 and Theorem 4.3. The first of these has the advantage that it does not depend on any order function and thus applies to a wider class of universes of separations; on the other hand, for those universes that do have an order function, the latter theorem is much more flexible and powerful, since it not only distinguishes all distinguishable profiles across all orders simultaneously, but also does so efficiently.

Our aim in this section is to establish Theorem 1.4 which combines the advantages of both Theorem 1.3 and Theorem 4.3 (without canonicity), i.e., which is not dependent on the existence of some order function, but which is as powerful and efficient as Theorem 4.3 if such an order function does exist.

Concretely, we shall answer the following question, which inspired our main theorem:

If $S_1 \subseteq S_2 \subseteq \dots \subseteq S_n$ is an ascending sequence of structurally submodular separations systems exhausting a universe of separations U , does there exist a nested set of separations which efficiently distinguishes all the maximal profiles in U ?

Let us substantiate this question with rigorous definitions of the terms involved.

We call a sequence $S_1 \subseteq S_2 \subseteq \dots \subseteq S_n \subseteq U$ of submodular separation systems in a universe U *compatible* if for all pairs $s_i \in S_i$ and $s_j \in S_j$ with $i \leq j$, either S_i contains at least two corner separations of s_i and s_j , or S_j contains at least three corner separations of s_i and s_j .

Observe that if U comes with a submodular order function $||$ and the S_i are defined as in Section 4.3, i.e., if S_i is the set of all separations in U of order $< i$, then the sequence $S_1 \subseteq S_2 \subseteq \dots \subseteq S_n \subseteq U$ is a compatible sequence of submodular separation systems.

A *profile* in $\mathcal{S} = (S_1, \dots, S_n)$ is a profile of one of the S_i .

A separation $s \in S_n$ *distinguishes* two profiles P and Q in \mathcal{S} if there are orientations of s such that $\vec{s} \in P$ and $\overleftarrow{s} \in Q$. The separation s distinguishes P and Q *efficiently* if $s \in S_i$ for every S_i which contains a separation that distinguishes P and Q .

Note once more that, as above, these notions of profiles and efficient distinguishers coincide with their usual definitions as given in Section 4.3 if U has a submodular order function and the S_i are the subsets of U containing all separations of order $< i$.

We also require a structural formulation of the concept of robustness from [9]:

A set \mathcal{P} of profiles in \mathcal{S} is *robust* if for all $P, Q, Q' \in \mathcal{P}$ the following holds: for every $\vec{r} \in Q \cap Q'$ with $\overleftarrow{r} \in P$ and every s which distinguishes Q and Q' efficiently, if $s \in S_j$, then there is an orientation \vec{s} of s such that either $(\overleftarrow{r} \vee \vec{s}) \in P$ or $(\vec{r} \vee \vec{s}) \in \vec{S}_j$.




Figure 1: Robustness.

With the above definitions we are now able to formally state and prove Theorem 1.4, which includes both Theorem 1.3 and the non-canonical Theorem 4.3 (and hence Theorem 1.1) as special cases:

Theorem 1.4. *If $\mathcal{S} = (S_1, \dots, S_n)$ is a compatible sequence of structurally submodular separation systems inside a universe U , and \mathcal{P} is a robust set of profiles in \mathcal{S} , then there is a nested set N of separations in U which efficiently distinguishes all the distinguishable profiles in \mathcal{P} .*

Since the proof of Theorem 1.4 runs along very similar lines as the proof of Theorem 4.3 in the previous section we only sketch it here:

Sketch of proof. For every pair P, P' of distinguishable profiles in \mathcal{P} let $A_{P,P'}$ be the set of all $s \in S_n$ that distinguish P and P' efficiently. The assertion of Theorem 1.4 follows directly from Theorem 1.5 if we can show that the family \mathcal{A} of these sets $A_{P,P'}$ splinters.

So let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$ be given. If r and s are nested there is nothing to show, so suppose they cross. Let i and j be minimal integers such that $r \in S_i$ and $s \in S_j$; we may assume without loss of generality that $i \leq j$.

If r distinguishes Q and Q' then $r \in A_{Q,Q'}$, so suppose not, that is, suppose that some orientation \vec{r} of r lies in both Q and Q' .

If one of the two corner separations $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \vec{s}$ lies in \vec{S}_j then that separation distinguishes Q and Q' by consistency and the profile property and hence would lie in $A_{Q,Q'}$. So we may suppose that neither of these two corner separations lies in \vec{S}_j . The compatibility of \mathcal{S} then implies that both of the other two corner separations, $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \vec{s}$, lie in S_i .

By possibly switching the roles of P and P' we may assume that $\vec{r} \in P'$ and $\vec{r} \in P$. Then the robustness of \mathcal{P} implies that P contains either $\vec{r} \vee \vec{s}$ or $\vec{r} \vee \vec{s}$. This corner separation then lies in $A_{P,P'}$ due to the consistency of P' . \square

Theorem 1.4 directly implies both Theorem 1.3 and the non-canonical Theorem 4.3: for the first theorem, consider the singleton sequence $S_1 = S$; and for the latter, take as S_i the set of all separations of order $< i$ and let n be large enough that $S_n = U$.

4.5 Clique-separations in finite graphs

For a finite graph G a separation (A, B) of G is a *clique separation* if the induced subgraph $G[A \cap B]$ is a complete graph. Clique separations in graphs have been studied by various people over the course of the last century [18, 22]. More recently clique separations have received quite some attention in theoretical computer science (see for instance [1, 5, 19]) following Tarjan's work [21] on their algorithmic aspects.

In [11] it was shown that the theory of submodular separation systems can be applied to clique separations of finite graphs to deduce the existence of certain nested distinguishing sets. Using Theorem 1.5 directly instead of Theorem 1.3, we are able to obtain a stronger result than the one given in [11], much in the same way that Theorem 1.4 improves upon Theorem 1.3.

For this section let $G = (V, E)$ be a finite graph, $\vec{U} = \vec{U}(G)$ the universe of separations of G , and $\vec{S} = \vec{S}(G) \subseteq \vec{U}$ the separation system of all clique separations of G . Consequently $\vec{S}_k = \vec{S}_k(G)$ is the set of all clique-separations in G of order less than k , i.e., the set of all $(A, B) \in \vec{S}$ such that $|A \cap B| < k$.

It was shown in [11, Lemma 17] that S is a submodular separation system. Following their proof, we can show that in fact every $S_k \subseteq S$ is a submodular separation system, and that these extend each other in a way similar to the ordinary S_k of G :

Lemma 4.5. *Let r and s be two crossing clique separations with $|r| \leq |s|$. Then there are orientations \vec{r} and \vec{s} of r and s such that $(\vec{r} \wedge \vec{s})$, $(\vec{r} \wedge \vec{s})$, and $(\vec{r} \wedge \vec{s})$ are clique separations with $|\vec{r} \wedge \vec{s}| \leq |r|$ and $|\vec{r} \wedge \vec{s}| \leq |r|$ as well as $|\vec{r} \wedge \vec{s}| \leq |s|$. Moreover, if $|\vec{r} \wedge \vec{s}| = |r| = |s|$, then $(\vec{r} \wedge \vec{s})$ is also a clique separation with $|\vec{r} \wedge \vec{s}| \leq |r|$.*

Proof. Let $s = \{A, B\}$ and $t = \{C, D\}$ be two crossing clique separations of G with $|r| \leq |s|$. Since $C \cap D$ is a separator of G , and all vertices in $A \cap B$ are pairwise adjacent,

$A \cap B$ must be a subset of either C or D . Similarly $C \cap D$ must be a subset of either A or B . By renaming the sets if necessary we may assume that $A \cap B \subseteq C$ and $C \cap D \subseteq A$. We orient r as $\vec{r} = (A, B)$ and s as $\vec{s} = (C, D)$; let us show that these orientations are as claimed.

Observe first that the separators of both $(\vec{r} \wedge \vec{s})$ and $(\vec{r} \wedge \vec{s})$ are subsets of $A \cap B$, showing that these are clique separations of order at most $|r| = |A \cap B|$. Similarly, the separator of the corner separation $(\vec{r} \wedge \vec{s})$ is a subset of $C \cap D$, and hence $(\vec{r} \wedge \vec{s})$ is a clique separation of order at most $|s| = |C \cap D|$.

Finally, suppose that $|\vec{r} \wedge \vec{s}| = |r| = |s|$. Then, since the separator of $(\vec{r} \wedge \vec{s})$ is a subset of both $A \cap B$ and of $C \cap D$, this separator must in fact be equal to both $A \cap B$ and $C \cap D$. Consequently the separator of $(\vec{r} \wedge \vec{s})$ also equals $A \cap B = C \cap D$, which shows that $(\vec{r} \wedge \vec{s})$ is a clique separation of order at most r . \square

We can now consider profiles in G with respect to these separation systems. We will use the same notion of profiles as was introduced in [9]: a *profile* P of order k is a consistent orientation of S_k satisfying the profile property

$$\forall \vec{r}, \vec{s} \in P: (\vec{r} \wedge \vec{s}) \notin P. \quad (\text{P})$$

Every hole in G (i.e. an induced cycle of length at least 4) defines a profile P of order $|V|$ in G by letting P contain a separation $(A, B) \in \vec{S}$ of order less than $|V|$ if and only if that hole is contained in $G[B]$. In an analogous way every clique of size k defines a profile of order k in G . Let us denote by \mathcal{P}_k the set of all profiles of order k .

As usual, given two distinguishable profiles P and P' , let

$$A_{P,P'} := \{a \in S \mid a \text{ distinguishes } P, P' \text{ efficiently}\}.$$

We will show that the collection of these $A_{P,P'}$ splinters.

Lemma 4.6. *For any set \mathcal{P} of profiles the collection $(A_{P,P'} \mid P, P' \text{ distinguishable profiles})$ splinters.*

Proof. Let P, P' and Q, Q' be two pairs of distinguishable profiles in \mathcal{P} and let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$ be two distinct separations. We need to show that we have either $r \in A_{Q,Q'}$ or $s \in A_{P,P'}$, or that some corner separation of r and s lies in $A_{P,P'} \cup A_{Q,Q'}$. If r and s are nested then the latter is immediate, so suppose that r and s cross. By switching their roles if necessary we may further assume that $|r| \leq |s|$.

Since Q orients s , and $|r| \leq |s|$, the profile Q contains some orientation \vec{r} of r . Similarly Q' contains some orientation of r . If $\vec{r} \in Q'$ then r distinguishes Q and Q' , and by $|r| \leq |s|$ it does so efficiently, giving $r \in A_{Q,Q'}$. So suppose that $\vec{r} \in Q$.

By Lemma 4.5 at least three of the corner separations of r and s are clique separations of order at most $|s|$. Thus at least one of $(\vec{r} \wedge \vec{s})$ and $(\vec{r} \wedge \vec{s})$ is a clique separation of order at most $|s|$. This corner separation then distinguishes Q and Q' by the profile property, and in fact it does so efficiently, since its order is at most $|s|$, yielding the desired corner separation in $A_{Q,Q'}$. \square

It is now straightforward to use Theorem 1.5 to obtain the following theorem:

Theorem 4.7. *There is a nested set of separations which efficiently distinguishes all the distinguishable profiles in $\bigcup_{i=1}^n \mathcal{P}_i$.*

Proof. By Lemma 4.6, we can apply Theorem 1.5 to $(A_{P,P'} \mid P, P'$ distinguishable profiles), resulting in the claimed nested subset. \square

In particular, for any two holes, a hole and a clique, or two cliques if there is a clique separation which distinguishes them, then our nested set contains one such separation of minimal order. As usual, such a nested set can be transformed into a tree-decomposition of G (see [8] for details). Thus G admits a tree-decomposition whose adhesion sets are cliques and which efficiently distinguishes all the holes and cliques distinguishable by clique separations in G . Such a decomposition is similar to, but not exactly the same as, the decomposition constructed by R. E. Tarjan in [21].

We will see in Section 6.2 that such a decomposition can in fact be chosen canonically, i.e., to be invariant under automorphisms of G .

5 Canonical Splinter Theorem

As we saw in the previous section, Theorem 1.5 is already strong enough to imply most of Theorem 4.3, but crucially does not guarantee the canonicity asserted in (iii). In this section we wish to prove a version of Theorem 1.5 using a stronger set of assumptions from which we can deduce Theorem 4.3 in full: we want to find, for a family $\mathcal{A} = (A_i \mid i \in I)$ of subsets of some universe U , a nested set $N = N(\mathcal{A})$ meeting all the A_i that is *canonical*, i.e., which only depends on invariants of \mathcal{A} . More formally, we want to find $N = N(\mathcal{A})$ in such a way that if $\mathcal{A}' = (A'_i \mid i \in I)$ is another family of subsets of some other universe U' that also meets the assumptions of our theorem, and φ is an isomorphism of separation systems between $\bigcup_{i \in I} A_i$ and $\bigcup_{i \in I} A'_i$ with $\varphi(A_i) = A'_i$ for all $i \in I$, we ask that $N(\mathcal{A}') = \varphi(N(\mathcal{A}))$. In particular, the nested set found by our theorem should not depend on the universe into which the family \mathcal{A} is embedded.

The assumptions of Theorem 1.5 are not sufficient to guarantee the existence of such a canonical set. Consider the example where we have just two separations, s and t , which are crossing and let $\mathcal{A} = (A_1) = (\{s, t\})$. Note that \mathcal{A} splinters, but there may be an automorphism that swaps the two separations so the choice of any single one of them is non-canonical. Since the separations are crossing we cannot use both of them for our nested set either.

For obtaining a canonical nested set, one crucial ingredient will be the notion of *extremal* elements of a set of separations, which was already used in [9]. Given a set $A \subseteq U$ of (unoriented) separations, an element $a \in A$ is *extremal* in A , or an *extremal element* of A , if a has some orientation \vec{a} that is a maximal element of \vec{A} . (Recall that \vec{A} is the set of orientations of separations in A .) The set of extremal elements of a set of separations is an invariant of separation systems in the following sense: if E is the set of extremal elements of some set $A \subseteq S$ of separations, and φ is an isomorphism

between \vec{S} and some other separation system, then $\varphi(E)$ is precisely the set of extremal separations of $\varphi(A)$. Moreover, the extremal separations of a set $A \subseteq U$ are nested with each other under relatively weak assumptions: for instance, it suffices that for any two separations in A at least two of their corner separations also lie in A .

Let us formally state a set of assumptions under which we can prove a canonical version of Theorem 1.5. Given two separations r and s and two of their corner separations c_1 and c_2 , we say that c_1 and c_2 are *from different sides of r* if, for orientations of c_1 , r , and s with $\vec{c}_1 = (\vec{r} \wedge \vec{s})$, there is an orientation \vec{c}_2 of c_2 such that either $\vec{c}_2 = (\vec{r} \wedge \vec{s})$ or $\vec{c}_2 = (\vec{r} \wedge \vec{s})$. Note that c_1 and c_2 being from different sides of r does not imply that c_1 and c_2 are distinct separations; consider for instance the edge case that $r = s = c_1 = c_2$.

Let $\mathcal{A} = (A_i \mid i \in I)$ be a finite collection of non-empty finite subsets of U and let \preceq be any partial order on I . We write $i \prec j$ if and only if $i \preceq j$ and $i \neq j$. We say that \mathcal{A} *splinters hierarchically* if for all $a_i \in A_i$ and $a_j \in A_j$ the following two conditions hold:

- (1) If $i \prec j$, either some corner separation of a_i and a_j lies in A_j , or two corner separations of a_i and a_j from different sides of a_i lie in A_i .
- (2) If neither $i \prec j$ nor $j \prec i$, there are $k \in \{i, j\}$ and corner separations c_1 and c_2 of a_i and a_j from different sides of a_k such that $c_1 \in A_k$ and $c_2 \in A_i \cup A_j$.




Figure 2: The possible configurations in (2) in the definition of *splinter hierarchically*, up to symmetry.

In particular if \preceq is the trivial partial order on I in which all $i \neq j$ are incomparable then \mathcal{A} splinters hierarchically if and only if (2) holds for all $a_i \in A_i$ and $a_j \in A_j$; this special case which ignores the partial order on I is perhaps the cleanest form of an assumption that suffices for a canonical nested set meeting all A_i in \mathcal{A} . The reason we need to allow a partial order \preceq on I and the slightly weaker condition in (1) for comparable elements of I is that otherwise we would not be able to deduce Theorem 4.3 in full from our main theorem of this section due to a quirk in the way that robustness is defined for profiles in [9] (see Section 6).

Our first lemma enables us to find a canonical nested set inside $\bigcup_{i \in I} A_i$ for a collection of sets A_i whose indexing set is an antichain:

Lemma 5.1. *Let $(A_i \mid i \in I)$ be a collection of subsets of U that splinters hierarchically. If $K \subseteq I$ is an antichain in \preceq , then the set of extremal elements of $\bigcup_{k \in K} A_k$ is nested.*

Proof. Suppose that $K \subseteq I$ is an antichain and that for some $i, j \in K$ there are $a_i \in A_i$ and $a_j \in A_j$ such that a_i and a_j are extremal in $\bigcup_{k \in K} A_k$ but cross. Let \vec{a}_i and \vec{a}_j be

the orientations of a_i and a_j witnessing their extremality. As a_i and a_j cross, there are three ways of orienting a_i and a_j such that the supremum of this orientation is strictly larger than \vec{a}_i or \vec{a}_j . Hence none of these corner separations can lie in $A_i \cup A_j$, since that would contradict the maximality of \vec{a}_i or \vec{a}_j in $\bigcup_{k \in K} \vec{A}_k$. On the other hand, since neither $i \prec j$ nor $j \prec i$, by condition (2) and the assumption that a_i and a_j cross there are at least two orientations of a_i and a_j whose corresponding supremum lies in $A_i \cup A_j$, causing a contradiction to the extremality of a_i and a_j . \square

We are now able to prove a canonical splinter theorem by repeatedly applying Lemma 5.1 to the collection of the A_i of \preceq -minimal index that have not yet been met by the nested set constructed so far:

Theorem 1.6 (Canonical splinter theorem). *Let U be a universe of separations and let $\mathcal{A} = (A_i \mid i \in I)$ be a collection of subsets of U that splinters hierarchically with respect to a partial order \preceq on I . Then there exists a nested set $N = N(\mathcal{A})$ meeting every A_i in \mathcal{A} .*

Moreover, $N(\mathcal{A})$ is canonical: if φ is an isomorphism of separation systems between $\bigcup_{i \in I} \vec{A}_i$ and a subset of some universe U' such that the family $\varphi(\mathcal{A}) := (\varphi(A_i) \mid i \in I)$ splinters hierarchically with respect to \preceq , then $N(\varphi(\mathcal{A})) = \varphi(N(\mathcal{A}))$.

Proof. We proceed by induction on $|I|$. If $|I| = 1$ we can choose as N the set of extremal elements of A_i , which is nested by Lemma 5.1 and clearly canonical.

So suppose that $|I| > 1$ and that the claim holds for all smaller index sets. Let K be the set of minimal elements of I with respect to \preceq . By Lemma 5.1 the set $E = E(\mathcal{A})$ of extremal elements of $\bigcup_{k \in K} A_k$ is nested. Let $J \subseteq I$ be the set of indices of all those A_j that do not meet E , and for $j \in J$ let A'_j be the set of all elements of A_j that are nested with E . We claim that the collection $\mathcal{A}' = (A'_j \mid j \in J)$ splinters hierarchically with respect to \preceq on J . This follows from Lemma 3.1 as soon as we show that each A'_j is non-empty.

To see that each A'_j is non-empty, for $j \in J$ let a_j be an element of A_j that crosses as few elements of E as possible. We wish to show that a_j is nested with E and thus $a_j \in A'_j$. So suppose that a_j crosses some separation in E , that is, some $a_i \in A_i \cap E$ with $i \in I \setminus J$. Since i is a minimal element of I we have either $i \preceq j$ or that i and j are incomparable. We shall treat these cases separately.

Consider first the case that $i \prec j$. By condition (1) of splintering hierarchically, either some corner separation of a_i and a_j lies in A_j , or two corner separations of a_i and a_j from different sides of a_i lie in A_i . The first of these possibilities contradicts the choice of a_j , since that corner separation in A_j would cross fewer elements of E by Lemma 3.1. On the other hand, the latter of these possibilities contradicts the choice of a_i as an extremal element of $\bigcup_{k \in K} A_k$. Thus the case $i \preceq j$ is impossible.

Let us now consider the case that i and j are incomparable. Again, by the choice of a_j , none of the corner separations of a_i and a_j can lie in A_j by Lemma 3.1. Therefore condition (2) of splintering hierarchically yields the existence of a corner separation of a_i and a_j in A_i for each side of a_i ; this, however, contradicts the extremality of a_i in $\bigcup_{k \in K} A_k$ as before.

Therefore each of the sets A'_j with $j \in J$ is non-empty, and hence the collection $\mathcal{A}' = (A'_j \mid j \in J)$ splinters hierarchically with respect to \preceq . Since $|J| < |I|$ we may apply the induction hypothesis to this collection to obtain a canonical nested set $N' = N(\mathcal{A}')$ meeting all A'_j . Now $N = N' \cup E$ is a nested subset of U which meets every A_i for $i \in I$. It remains to show that N is canonical.

To see that N is canonical let φ be an isomorphism of separation systems between $\bigcup_{i \in I} \vec{A}_i$ and a subset of some universe U' such that $\varphi(\mathcal{A})$ splinters hierarchically with respect to \preceq in U' . Then $\varphi(E) = E(\varphi(\mathcal{A}))$, i.e., the set of extremal elements of $\bigcup_{i \in I} \varphi(A_i)$ is exactly $\varphi(E)$. Therefore $\varphi(E)$ meets $\varphi(A_i)$ if and only if E meets A_i . Consequently the restriction of φ to $\bigcup_{j \in J} \vec{A}_j$ is an isomorphism of separation systems between $\bigcup_{j \in J} \vec{A}_j$ and its image in U' with the property that $\varphi(\mathcal{A}')$ splinters hierarchically with respect to \preceq on J . Moreover, for $j \in J$, the image $\varphi(A'_j)$ of A'_j is exactly the set of those separations in $\varphi(A_j)$ that are nested with $\varphi(E)$.

Thus we can apply the induction hypothesis to find that $N(\varphi(\mathcal{A}')) = \varphi(N(\mathcal{A}'))$. Together with the above observation that $\varphi(E(\mathcal{A})) = E(\varphi(\mathcal{A}))$ this gives

$$\varphi(N(\mathcal{A})) = \varphi(E(\mathcal{A})) \cup \varphi(N(\mathcal{A}')) = E(\varphi(\mathcal{A})) \cup N(\varphi(\mathcal{A}')) = N(\varphi(\mathcal{A})),$$

concluding the proof. \square

6 Applications of the Canonical Splinter Theorem

In this section we apply Theorem 1.6 to obtain a short proof of Theorem 4.3, to strengthen Theorem 4.7 for clique separations so as to make it canonical, and finally to establish a canonical tree-of-tangles theorem for another type of separations, so-called circle separations.

6.1 Robust profiles

Having established Theorem 1.6 in the previous section, we are now ready to derive the full version of Theorem 4.3. For this let $U = (\vec{U}, \leq, *, \vee, \wedge, | \cdot |)$ be a submodular universe of separations and \mathcal{P} a robust set of profiles in U , and let I be the set of all pairs of distinguishable profiles in \mathcal{P} . As in Section 4.3, for $\{P, P'\} \in I$ we let

$$A_{P, P'} := \{a \in U \mid a \text{ distinguishes } P \text{ and } P' \text{ efficiently}\},$$

and let $\mathcal{A}_{\mathcal{P}}$ be the family $(A_{P, P'} \mid \{P, P'\} \in I)$. We furthermore define a partial order \preceq on I by letting $\{P, P'\} \prec \{Q, Q'\}$ if and only if the order of some element of $A_{P, P'}$ is strictly lower than the order of some element of $A_{Q, Q'}$. Note that the separations in a fixed $A_{P, P'}$ all have the same order.

We shall be able to deduce Theorem 4.3 from Theorem 1.6 as soon as we show that $\mathcal{A}_{\mathcal{P}}$ splinters hierarchically.

Lemma 6.1. $\mathcal{A}_{\mathcal{P}}$ splinters hierarchically with respect to \preceq .

Proof. Let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$ be given. By switching their roles if necessary we may assume that $|r| \leq |s|$. Then Q and Q' both orient r ; we may assume without loss of generality that $\vec{r} \in Q$. We will make a case distinction depending on the way Q' orients r .

Let us first treat the case that Q and Q' orient r differently, i.e., that $\vec{r} \in Q'$. Then r distinguishes Q and Q' and hence $|r| = |s|$ by the efficiency of s . This implies that $\{P, P'\}$ and $\{Q, Q'\}$ are either the same pair or else incomparable in \preceq . We may assume further without loss of generality that $\vec{s} \in Q$ and $\vec{s} \in Q'$. Consider now the two corner separations $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$: if at least one of these two has order at most $|s|$, then this corner separation would distinguish Q and Q' by the profile property. The efficiency of s would then imply that this corner separation has order exactly $|s|$ and hence lies in $A_{Q,Q'}$. The submodularity of the order function implies that this is the case for at least one, and therefore for both of these corner separations, yielding the existence of two corner separations of r and s from different sides of s in $A_{Q,Q'}$ and showing that (2) is satisfied.

Let us now consider the case that Q and Q' orient r in the same way, i.e., that $\vec{r} \in Q'$. We make a further split depending on whether $|r| = |s|$ or $|r| < |s|$.

Suppose first that $|r| = |s|$; then neither $\{P, P'\} \prec \{Q, Q'\}$ nor $\{Q, Q'\} \prec \{P, P'\}$. We may assume that P and P' orient s in the same way: for if P and P' orient s differently, we may switch the roles of r and s as well as $\{P, P'\}$ and $\{Q, Q'\}$ and apply the above case. So suppose that both of P and P' contain \vec{s} , say. Then neither of the corner separations $\vec{r} \vee \vec{s}$ nor $\vec{r} \wedge \vec{s}$ can have order strictly less than $|r| = |s|$, as these corner separations would distinguish Q and Q' or P and P' , respectively, and would therefore contradict the efficiency of s or of r , respectively. The submodularity of $|\cdot|$ now implies that both of these corner separations have order exactly $|r| = |s|$ and hence lie in $A_{Q,Q'}$ and $A_{P,P'}$, respectively, showing that (2) holds.

Finally, let us suppose that $|r| < |s|$; then $\{P, P'\} \prec \{Q, Q'\}$. Consider the two corner separations $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$: if both of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ have order strictly greater than $|s|$, then by the submodularity of the order function both of the other two corner separations $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ have order strictly smaller than $|r|$. By the robustness of \mathcal{P} one of these two corner separations would distinguish P and P' , contradicting the efficiency of r .

Thus we may assume at least one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ has order at most $|s|$. Then that corner separation distinguishes Q and Q' . In fact, it does so efficiently and hence lies in $A_{Q,Q'}$, showing that (1) holds and concluding the proof. \square

We are now ready to deduce the full Theorem 4.3 from Theorem 1.6:

Theorem 4.3 (Canonical tangle-tree theorem for separation universes [9, Theorem 3.6]). *Let $U = (\vec{U}, \preceq, *, \vee, \wedge, |\cdot|)$ be a submodular universe of separations. Then for every robust set \mathcal{P} of profiles in U there is a nested set $T = T(\mathcal{P}) \subseteq U$ of separations such that:*

- (i) every two profiles in \mathcal{P} are efficiently distinguished by some separation in T ;
- (ii) every separation in T efficiently distinguishes a pair of profiles in \mathcal{P} ;

(iii) for every automorphism α of \vec{U} we have $T(\mathcal{P}^\alpha) = T(\mathcal{P})^\alpha$; (canonicity)

(iv) if all the profiles in \mathcal{P} are regular, then T is a regular tree set.

Proof. By Lemma 6.1 the family $\mathcal{A}_\mathcal{P}$ splinters hierarchically. Thus we can apply Theorem 1.6 to $\mathcal{A}_\mathcal{P}$ to obtain a nested set $N = N(\mathcal{A}_\mathcal{P})$ which meets every $A_{P,P'}$. Clearly, N satisfies (i), (ii) and (iv) of Theorem 4.3.

To see that N satisfies (iii), let α be an automorphism of \vec{U} . Then the restriction of α to $\bigcup_{\{P,P'\} \in I} \vec{A}_{P,P'}$ is an isomorphism of separation systems onto its image in \vec{U} . We therefore have, by Theorem 1.6, that $\alpha(N(\mathcal{A}_\mathcal{P})) = N(\alpha(\mathcal{A}_\mathcal{P}))$. For every $A_{P,P'}$ in $\mathcal{A}_\mathcal{P}$ we have that $\alpha(A_{P,P'})$ is precisely the set of those separations in U which distinguish P^α and P'^α efficiently; in other words, we have $\alpha(\mathcal{A}_\mathcal{P}) = \mathcal{A}_{\mathcal{P}^\alpha}$, showing that (iii) is satisfied. \square

6.2 Clique separations

Regarding the profiles of clique separations discussed in Section 4.5, Lemma 4.5 not only suffices to show that the sets $A_{P,P'}$ splinters, but can be used to show that the collection of these $A_{P,P'}$ even splinters hierarchically, allowing us to apply Theorem 1.6: for this we simply define the same partial order \preceq on the set of pairs $\{P, P'\}$ as in the previous section, that is, $\{P, P'\} \prec \{Q, Q'\}$ if and only if $|r| < |s|$ for some (equivalently: for all) $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$.

To see this, let P, P' and Q, Q' be distinguishable pairs of profiles of clique separations. Let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$, and suppose without loss of generality that $|r| \leq |s|$. If r and s are nested, then r and s themselves are corner separations of r and s that lie in $A_{P,P'}$ and $A_{Q,Q'}$, respectively. However, if r and s cross, then by Lemma 4.5 there are orientations of r and s such that $|\vec{r} \wedge \vec{s}|, |\vec{r} \wedge \vec{s}'| \leq |r|$ and $|\vec{r} \wedge \vec{s}|, |\vec{r} \wedge \vec{s}'|, |\vec{r}' \wedge \vec{s}| \leq |s|$. By switching their roles if necessary we may assume that $\vec{r}' \in P$ and $\vec{r} \in P'$, and likewise that $\vec{s} \in Q$ and $\vec{s}' \in Q'$.

Since $(\vec{r}' \wedge \vec{s}), (\vec{r} \wedge \vec{s}) \leq \vec{s}$ and $\vec{s} \in Q'$, the profile Q' contains both of these corner separations by consistency. On the other hand, by the assumption that $|r| \leq |s|$, the separation r gets oriented by Q , and consequently by the profile property Q must contain the inverse of one of those two corner separations. This corner separation then distinguishes Q and Q' , and in fact it does so efficiently, since its order is at most $|s|$, meaning that this corner separation lies in $A_{Q,Q'}$. Therefore, if $|r| < |s|$, condition (1) of splintering hierarchically is satisfied.

So suppose further that $|r| = |s|$, and let us check that (2) of splintering hierarchically is satisfied. Observe that, similarly as above, P orients s , and P' contains both $(\vec{r}' \wedge \vec{s})$ and $(\vec{r} \wedge \vec{s})$ by consistency with $\vec{r} \in P'$, implying as before that one of $(\vec{r} \wedge \vec{s})$ and $(\vec{r}' \wedge \vec{s})$ also efficiently distinguishes P and P' , i.e., is an element of $A_{P,P'}$. If this corner separation in $A_{P,P'}$ and the corner separation in $A_{Q,Q'}$ found above are from different sides of either r or s , then (2) of splintering hierarchically would be satisfied. So suppose not; that is, suppose that $(\vec{r} \wedge \vec{s})$ distinguishes both P and P' as well as Q and Q' efficiently. In particular $|\vec{r} \wedge \vec{s}| = |r| = |s|$, and hence by the last part of Lemma 4.5, all four corner separations of r and s have order at most $|r|$. Consequently, since P' orients

s , one of $(\vec{r} \wedge \vec{s})$ and $(\vec{r} \wedge \vec{s})$ distinguishes P and P' efficiently, which one depending on whether $\vec{s} \in P'$ or $\vec{s} \in P'$. In either case we have found a corner separation of r and s in $A_{P,P'}$, which together with $(\vec{r} \wedge \vec{s}) \in A_{Q,Q'}$ witnesses that (2) is fulfilled.

Therefore, by Theorem 1.6 we get that we can choose the set in Theorem 4.7 canonically:

Theorem 6.2. *For every set \mathcal{P} of profiles of clique separations of a graph G , there is a nested set $N = N(\mathcal{P})$ of separations which efficiently distinguishes all the distinguishable profiles in \mathcal{P} and is canonical, that is, such that $N(\mathcal{P}^\alpha) = N(\mathcal{P})^\alpha$ for every automorphism α of the underlying graph G .*

Proof. Every automorphism of G induces an automorphism of the separation system. Hence we can obtain the claimed nested set by applying Theorem 1.6 to the family of the sets $A_{P,P'}$ of those clique separations which efficiently distinguish the pair P, P' of distinguishable profiles in \mathcal{P} . \square

6.3 Circle separations

Another special case of separation systems are those of *circle separations* discussed in [11]: given a fixed cyclic order on a ground-set V , a *circle separation* of V is a bipartition (A, B) of V into two disjoint intervals in the cyclic order. Observe that the set of all circle separations is not closed under joins and meets and hence not a sub-universe of the universe of all bipartitions of V :

Example 6.3. Consider the natural cyclic order on the set $V = \{1, 2, 3, 4\}$. The bipartitions $(\{1\}, \{2, 3, 4\})$ and $(\{3\}, \{4, 1, 2\})$ of V are circle separations. However, their supremum in the universe of all bipartitions of V is $(\{1, 3\}, \{2, 4\})$, which is not a circle separation.

Let V be a ground-set with a fixed cyclic order and $U = (\vec{U}, \leq, *, \vee, \wedge, | |)$ the universe of all bipartitions of V with a submodular order function $| |$. Let $S \subseteq U$ be the set of all separations in U that are circle separations of V . Consequently we denote by S_k the set of all those circle separations in S whose order is $< k$.

Given fixed integers $m \geq 1$ and $n > 3$, we call a consistent orientation of S_k a *k-tangle* in S if it has no subset in

$$\mathcal{F} = \mathcal{F}_m^n := \left\{ F \subseteq 2^{\vec{U}} \mid |\cap_{(A,B) \in F} B| < m \text{ and } |F| < n \right\}.$$

A *tangle* in S is then a *k-tangle* for some k , and a *maximal tangle* in S is a tangle not contained in any other tangle in S . As usual, two tangles are *distinguishable* if neither of them is a subset of the other; a separation s *distinguishes* two tangles if they orient s differently, and s does so *efficiently* if it is of minimal order among all separations in S distinguishing that pair of tangles.

Using Theorem 1.6 we can show that there is a canonical nested set of circle separations which efficiently distinguishes all distinguishable tangles in S :

Theorem 6.4. *The set S of all circle separations of V contains a tree set $T = T(S)$ that efficiently distinguishes all distinguishable tangles of S . Moreover, this tree set T can be chosen canonically, i.e., so that for every automorphism α of S we have $T(S^\alpha) = T(S)^\alpha$.*

In order to prove Theorem 6.4 we need the following short lemma:

Lemma 6.5. *Let r and s be two circle separations of V . If r and s cross then all four corner separations of r and s are again circle separations.*

Proof. Let $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$. Since r and s cross, the sets $A \cap C$ and $B \cap D$ are non-empty and moreover intervals in the cyclic order. Thus $B \cup D$ is also an interval and therefore $\vec{r} \wedge \vec{s} = (A \cap C, B \cup D)$ is indeed a circle separation. \square

Let us now prove Theorem 6.4.

Proof of Theorem 6.4. For every pair P, P' of distinguishable tangles in S let $A_{P, P'}$ be the set of all circle separations that efficiently distinguish P and P' . We define a partial order \preceq on the set of all pairs of distinguishable tangles by letting $\{P, P'\} \prec \{Q, Q'\}$ for two distinct such pairs if and only if the separations in $A_{P, P'}$ have strictly lower order than those in $A_{Q, Q'}$.

Let us show that the collection of these sets $A_{P, P'}$ splinters hierarchically; the claim will then follow from Theorem 1.6.

For this let P, P' and Q, Q' be two distinguishable pairs of tangles in S and let $r \in A_{P, P'}$ and $s \in A_{Q, Q'}$. If r and s are nested, then r and s themselves are corner separations from different sides of r and s that lie in $A_{P, P'}$ and $A_{Q, Q'}$, respectively, in which case there is nothing to show.

So suppose that r and s cross. Then by Lemma 6.5 all corner separations of r and s are circle separations. By switching their roles if necessary we may assume that $|r| \leq |s|$; we shall treat the cases of $|r| < |s|$ and $|r| = |s|$ separately.

Let us first consider the case that $|r| < |s|$. Then $\{P, P'\} \prec \{Q, Q'\}$, so it suffices to show that (1) is satisfied, i.e., to find a corner separation of r and s in $A_{Q, Q'}$. Since Q and Q' both orient s , which is of higher order than r , both Q and Q' also orient r . By $|r| < |s|$ and the efficiency of s , r cannot distinguish Q and Q' . Thus some orientation \vec{r} of r lies in both Q and Q' .

By renaming them if necessary we may assume that $\vec{r} \in P$ and $\vec{r} \in P'$. Suppose now that one of $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \vec{s}$ has order at most $|s|$. Then Q and Q' would both orient that corner separation, and they would do so differently by the definition of a tangle. Thus that corner separation would lie in $A_{Q, Q'}$, as desired.

Hence we may assume that both of $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \vec{s}$ have order higher than $|s|$. Then, by submodularity, both $\vec{r} \wedge \vec{s}$ and $\vec{r} \wedge \vec{s}$ have order less than $|r|$. Therefore both of these corner separations get oriented by P and P' , but neither of them can distinguish P and P' by the efficiency of r . In fact by the consistency of P and P' we must have $(\vec{r} \wedge \vec{s}), (\vec{r} \wedge \vec{s}) \in P \cap P'$. However the set $\{\vec{r}, (\vec{r} \wedge \vec{s}), (\vec{r} \wedge \vec{s})\}$ lies in \mathcal{F} , contradicting the assumption that P and P' are tangles in S .

It remains to deal with the case that $|r| = |s|$ and show that (2) is satisfied. For this we shall find corner separations from different sides of r or of s that lie in $A_{P, P'}$ and

$A_{Q,Q'}$, respectively. By the submodularity of the order function, and by switching the roles of r and s if necessary, we may assume that there are orientations of r and s such that both $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \vec{s}$ have order at most $|r|$. By possibly renaming \vec{s} and \vec{s} we may further assume that $\vec{r} \vee \vec{s}$ distinguishes P and P' . Then, by the efficiency of r , we must have $|\vec{r} \vee \vec{s}| = |r|$, and hence $|\vec{r} \vee \vec{s}| \leq |s|$ by submodularity. Recall that we assumed $|\vec{r} \vee \vec{s}| = |r| = |s|$, so one of $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \vec{s}$ must distinguish Q and Q' . Again, that corner separation must in fact distinguish Q and Q' efficiently, i.e., lie in $A_{Q,Q'}$. Now this corner separation together with $\vec{r} \vee \vec{s}$ witnesses that (2) holds. \square

References

- [1] A. Berry, R. Pogorelcnik, and G. Simonet, *An introduction to clique minimal separator decomposition*, Algorithms (Basel) **3** (2010), no. 2, 197–215, doi:[10.3390/a3020197](https://doi.org/10.3390/a3020197).
- [2] N. Bowler and J. Kneip, *Separations of sets* (2018), available at arXiv:[1809.00625](https://arxiv.org/abs/1809.00625).
- [3] J. Carmesin, *A short proof that every finite graph has a tree-decomposition displaying its tangles*, European J. Combin. **58** (2016), 61–65, doi:[10.1016/j.ejc.2016.04.007](https://doi.org/10.1016/j.ejc.2016.04.007).
- [4] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein, *Connectivity and tree structure in finite graphs*, Combinatorica **34** (2014), no. 1, 11–45, doi:[10.1007/s00493-014-2898-5](https://doi.org/10.1007/s00493-014-2898-5).
- [5] D. Coudert and G. Ducoffe, *Clique-decomposition revisited*, INRIA Sophia Antipolis - I3S, 2016, available at [hal-01266147](https://hal.inria.fr/hal-01266147).
- [6] R. Diestel, *Graph theory*, 5th ed., Graduate Texts in Mathematics, vol. 173, Springer, Berlin, 2017.
- [7] ———, *Abstract separation systems*, Order **35** (2018), no. 1, 157–170, doi:[10.1007/s11083-017-9424-5](https://doi.org/10.1007/s11083-017-9424-5).
- [8] ———, *Tree sets*, Order **35** (2018), no. 1, 171–192, doi:[10.1007/s11083-017-9425-4](https://doi.org/10.1007/s11083-017-9425-4).
- [9] R. Diestel, F. Hundertmark, and S. Lemanczyk, *Profiles of separations: in graphs, matroids, and beyond*, Combinatorica **39** (2019), no. 1, 37–75, doi:[10.1007/s00493-017-3595-y](https://doi.org/10.1007/s00493-017-3595-y).
- [10] R. Diestel, Ph. Eberenz, and J. Erde, *Duality theorems for blocks and tangles in graphs*, SIAM J. Discrete Math. **31** (2017), no. 3, 1514–1528, doi:[10.1137/16M1077763](https://doi.org/10.1137/16M1077763).
- [11] R. Diestel, J. Erde, and D. Weißauer, *Structural submodularity and tangles in abstract separation systems*, J. Combin. Theory Ser. A **167** (2019), 155–180, doi:[10.1016/j.jcta.2019.05.001](https://doi.org/10.1016/j.jcta.2019.05.001).
- [12] R. Diestel and J. Kneip, *Profinite separation systems*, Order (to appear), available at arXiv:[1804.01921](https://arxiv.org/abs/1804.01921).
- [13] R. Diestel and S. Oum, *Tangle-tree duality in abstract separation systems* (2017), available at arXiv:[1701.02509](https://arxiv.org/abs/1701.02509).
- [14] ———, *Tangle-Tree Duality: In Graphs, Matroids And Beyond*, Combinatorica (2019), doi:[10.1007/s00493-019-3798-5](https://doi.org/10.1007/s00493-019-3798-5).
- [15] R. Diestel and G. Whittle, *Tangles and the Mona Lisa* (2016), available at arXiv:[1603.06652](https://arxiv.org/abs/1603.06652).
- [16] Chr. Elbracht, J. Kneip, and M. Teegen, *Trees of tangles in infinite abstract separation systems*, in preparation.
- [17] P. Gollin and J. Kneip, *Representations of infinite tree sets* (2019), available at arXiv:[1908.10327](https://arxiv.org/abs/1908.10327).

- [18] R. Halin, *Über simpliziale Zerfällungen beliebiger (endlicher oder unendlicher) Graphen*, Math. Ann. **156** (1964), no. 3, 216–225, doi:[10.1007/BF01363288](https://doi.org/10.1007/BF01363288).
- [19] H.-G. Leimer, *Optimal decomposition by clique separators*, Discrete Math. **113** (1993), no. 1-3, 99–123, doi:[10.1016/0012-365X\(93\)90510-Z](https://doi.org/10.1016/0012-365X(93)90510-Z).
- [20] N. Robertson and P. D. Seymour, *Graph minors. X. Obstructions to tree-decomposition*, J. Combin. Theory Ser. B **52** (1991), no. 2, 153–190, doi:[10.1016/0095-8956\(91\)90061-N](https://doi.org/10.1016/0095-8956(91)90061-N).
- [21] R. E. Tarjan, *Decomposition by clique separators*, Discrete Math. **55** (1985), no. 2, 221–232, doi:[10.1016/0012-365X\(85\)90051-2](https://doi.org/10.1016/0012-365X(85)90051-2).
- [22] K. Wagner, *Über eine Eigenschaft der ebenen Komplexe*, Math. Ann. **114** (1937), 570–590, doi:[10.1007/BF01594196](https://doi.org/10.1007/BF01594196).

Christian Elbracht, Universität Hamburg, Hamburg, Germany

christian.elbracht@uni-hamburg.de

Jakob Kneip, Universität Hamburg, Hamburg, Germany

jakob.kneip@studium.uni-hamburg.de

Maximilian Teegen, Universität Hamburg, Hamburg, Germany

maximilian.teegen@uni-hamburg.de