

Obtaining trees of tangles from tangle-tree duality

CHRISTIAN ELBRACHT,
JAKOB KNEIP, AND MAXIMILIAN TEEGEN

We demonstrate the versatility of the tangle-tree duality theorem for abstract separation systems [7] by using it to prove tree-of-tangles theorems. This approach allows us to strengthen some of the existing tree-of-tangles theorems by bounding the node degrees in them. We also present a slight strengthening and simplified proof of the duality theorem, which allows us to derive a tree-of-tangles theorem also for tangles of different orders.

AMS 2000 SUBJECT CLASSIFICATIONS: 05C40, 05C83, 05C05.

KEYWORDS AND PHRASES: Tangles, tree-decompositions, abstract separation systems, tangle-tree duality, tree of tangles.

1. Introduction

Over the last twenty years, tangles, originally developed by Robertson and Seymour as a tool in their monumental graph minor project [13], have evolved a lot. Originally defined specifically just for graphs, they have since been generalized not only to other combinatorial structures like matroids [11] but even into an abstract setting in which concrete separations are replaced by an abstract poset with just some simple properties that reflect those that separations typically have [2, 5, 7]. In all these settings, the general idea of tangles is to use them as a method to indirectly capture highly cohesive substructures of various kinds, by deciding for every low-order separation on which side of that separation the desired structure lies.

Already in the original work by Robertson and Seymour the theory of tangles has two major theorems: the tree-of-tangles theorem and the tangle-tree duality theorem. These two form the main pillars of tangle theory, and thereby of a central aspect of graph minor theory.

The first of these theorems allows one to distinguish all the tangles in a tree-like way, displaying their relative position in the underlying combinatorial structure. One of the most abstract variants of the tree-of-tangles theorem reads as follows:

arXiv: [2011.09758](https://arxiv.org/abs/2011.09758)

Theorem 1.1 ([4, Theorem 6]). *Let \vec{S} be a structurally submodular separation system and \mathcal{P} a set of profiles of S . Then \vec{S} contains a tree set that distinguishes \mathcal{P} .*

Profiles are the most general class of objects that one can think of as tangles.

The tangle-tree duality theorem, on the other hand, provides a tree-like dual object to tangles which, if no tangle exists, serves as a witness that there can be no tangle. In this paper we demonstrate the versatility of the most abstract version of this duality theorem: we deduce Theorem 1.1 and some of its variations from the tangle-tree duality theorem, reducing the two pillars of abstract tangle theory to a single pillar.

In order to use tangle-tree duality to deduce tree-of-tangles theorems like Theorem 1.1, we exploit the generality of the most abstract version of the tangle-tree duality theorem, which reads as follows:

Theorem 1.2 (Tangle-tree duality theorem [7, Theorem 4.3]). *Let U be a universe of separations containing a finite separation system $S \subseteq U$ and let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars such that \mathcal{F} is standard for \vec{S} and \vec{S} is \mathcal{F} -separable. Then exactly one of the following statements holds:*

- *there is an \mathcal{F} -tangle of S ;*
- *there is an S -tree over \mathcal{F} .*

The strength of Theorem 1.2 lies in the flexibility it allows in the choice of \mathcal{F} . This set \mathcal{F} can be tailored to capture a wide variety of tangles and clusters, allowing Theorem 1.2 to be employed in a multitude of different settings ([6, 4]). The freedom in choosing and manipulating \mathcal{F} will also allow us to achieve our goal of deducing tree-of-tangles theorems from Theorem 1.2: by a clever choice of \mathcal{F} we can ensure that there is no \mathcal{F} -tangle of S , and that the S -tree over \mathcal{F} one then obtains will be a tree of tangles. We present multiple variations of this idea throughout this paper.

In terms of simplicity and brevity, reducing the tree-of-tangles theorem to the tangle-tree duality theorem in this way cannot compete with its direct proofs in [5], [4] or [10], our general purpose solution to obtaining tree-of-tangles theorems in a wide range of structures. (There, we showed an even more general theorem than Theorem 1.1 which no longer mentions tangles or profiles at all, but just talks about sets of separations fulfilling one simple-to-check condition.)

Instead of competing in terms of simplicity and brevity just for a proof of the tree-of-tangles theorem, the aim of this paper is to bridge the two parts of the theory needed for their classical proofs. This can be viewed in

two ways. Firstly, that we introduce tools from tangle-tree duality into the world of trees of tangles, which gives us a new method for building trees in this context very unlike the proofs in [5, 4, 10].

Secondly, and perhaps more importantly, from the perspective of tangle-tree duality this may be viewed as introducing a new range of ways of how to apply the duality theorem by a careful choice of \mathcal{F} . Previous applications of Theorem 1.2 all worked with largely similar choices of \mathcal{F} , all designed to capture some notion of ‘width’, whereas we specifically construct \mathcal{F} in such a way that no \mathcal{F} -tangle can exist, thereby making sure that Theorem 1.2 gives us the dual object which will be the desired tree-of-tangles.

A new result that we get from this method is that it allows us to bound the degrees of the nodes in a tree of tangles in some contexts. Getting such a degree condition out of the original proofs does not appear to be simple.

To achieve our last result, we prove a strengthened version of Theorem 1.2, which we present in Section 7.1 along with a simpler proof than the original one in [7].

The structure of this paper is as follows. In Section 2 we will repeat the required definitions from [5, 4, 2, 7, 6]. In Section 3 we prove our first basic tree-of-tangles theorem, for structurally submodular separation systems. A refined version of this argument will be given in Section 4, where we show that the approach via tangle-tree duality yields a bound on the degrees of the nodes in a tree of tangles. In Section 5 we present a more involved argument to obtain a tree of tangles that distinguishes a set of profiles *efficiently*. Again, this approach can be used to obtain a result about the degrees in such a tree, and we do so in Section 6. In Section 7 we prove a tree-of-tangles theorem for tangles of different orders. For that we need our stronger version of the tangle-tree duality theorem, which we state in Section 7.1. The proof of this stronger duality theorem in Section 7.1 also offers a new, and maybe simpler, proof of the original tangle-tree duality theorem Theorem 1.2. In our final section, Section 7.2, we then use this stronger tangle-tree duality theorem to obtain a tree-of-tangles theorem for profiles of different order.

2. Terminology and background

Since we combine the theory of tangle-tree duality and of trees of tangles we need the terminology of both. Consequently, this results in a large number of definitions which need to be understood for the comprehension of this paper. We employ the frameworks of [2, 4, 5, 6, 7]. For reference, we offer a recap of the definitions that we will use, split up according to their context. The cited sources provide more in-depth motivation of the respective set-ups.

2.1. Basic definitions of abstract separation systems (cf. [2])

A separation system $S = (\vec{S}, \leq, *)$ consists of a finite poset \vec{S} together with an involution $*$ which is order-reversing, i.e., $\vec{s} \leq \vec{t} \Leftrightarrow \vec{s}^* \geq \vec{t}^*$. We call the elements of \vec{S} (*oriented*) *separations* and denote, given an oriented separation $\vec{s} \in \vec{S}$, the image of \vec{s} under $*$ as the *inverse* $\vec{\bar{s}}$ of \vec{s} .

The pair $\{\vec{s}, \vec{\bar{s}}\}$ of \vec{s} together with its inverse $\vec{\bar{s}}$ is denoted as s and called the *underlying unoriented separation* of \vec{s} . Given s , we say that \vec{s} and $\vec{\bar{s}}$ are the *orientations* of s . The set of all the underlying unoriented separations for a set $\vec{T} \subseteq \vec{S}$ is denoted as T , so S is the set of all unoriented separations of separations in \vec{S} . Conversely, when given a set T of unoriented separations, we denote as \vec{T} the set of all orientations of separations in T . For brevity, we mean by the term ‘separations’ both oriented and unoriented separations if the intended meaning is clear from the context.

We say that $\vec{s} \in \vec{S}$ is *small* if it is less than its inverse, that is if $\vec{s} \leq \vec{\bar{s}}$. If there is an unoriented separation $r \neq s$ such that $\vec{s} \leq \vec{r}$ and $\vec{s} \leq \vec{\bar{r}}$, then \vec{s} is called *trivial*. Note that every trivial separation is small, since $\vec{s} \leq \vec{\bar{r}}$ implies that $\vec{\bar{s}} \geq \vec{r}$ and thus $\vec{s} \leq \vec{r} \leq \vec{\bar{s}}$.

The inverse $\vec{\bar{s}}$ of a small separation \vec{s} is called *cosmall* and likewise the inverse of a trivial separation is called *cotrivial*. A set of oriented separations is *regular* if it contains no cosmall separation.

We say that a separation s (and likewise its orientations) is *degenerate* if $\vec{s} = \vec{\bar{s}}$.

We say that two unoriented separation r and s from S are *nested* if they have orientations \vec{r}, \vec{s} such that $\vec{r} \leq \vec{s}$. Two separations r and s *cross* if they are not nested. Two oriented separations \vec{r} and \vec{s} cross or are nested if the underlying unoriented separations r and s cross or are nested, respectively. Note than in particular \vec{r} and \vec{s} can be nested even if they are incomparable, for instance if $\vec{r} \leq \vec{\bar{s}}$. A set of separations is called *nested* if its elements are pairwise nested.

A nested set T of unoriented separations is called a *tree set* in a separation system S , if T does not contain any separation s which has a trivial orientation in \vec{T} . A tree-set T is *regular* if \vec{T} is regular.

A universe of separations $U = (\vec{U}, \leq, *, \vee, \wedge)$ is a separation system $(\vec{U}, \leq, *)$ together with join and meet operators \vee, \wedge which turn the poset (\vec{U}, \leq) into a lattice. For universes of separations DeMorgan’s law holds:

$$(\vec{s} \vee \vec{t})^* = \vec{\bar{s}} \wedge \vec{\bar{t}}$$

Given two unoriented separations s and t in U , we call the unoriented separations corresponding to $\vec{s} \vee \vec{t}, \vec{s} \vee \vec{\bar{t}}, \vec{\bar{s}} \vee \vec{t}$ and $\vec{\bar{s}} \vee \vec{\bar{t}}$ the *corner separations*, or *corners* for short, of s and t .

One often-used property of separation universes is the so-called *fish lemma*:

Lemma 2.1 ([2, Lemma 3.2]). *Let $r, s \in U$ 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 .*

Given a separation system S , a subset $O \subseteq \vec{S}$ is *antisymmetric* if it contains at most one orientation of every separation, i.e., $|O \cap \{\vec{s}, \vec{\bar{s}}\}| \leq 1$ for every $s \in S$.

An *orientation* of S is an antisymmetric subset $O \subseteq \vec{S}$ such that $\vec{s} \in O$ or $\vec{\bar{s}} \in O$ for every $s \in S$. Such an orientation O is *consistent* if O does not contain any \vec{r} and \vec{s} such that $\vec{r} \leq \vec{s}$ and $r \neq s$.

Some universes U come with an *order function*, a function $|\cdot|: \vec{U} \rightarrow \mathbb{N}_0$ which is invariant under $*$, that is, for any $s \in U$, we have $|\vec{s}| = |\vec{\bar{s}}| = |s|$. Such an order function is called *submodular* if, for all $\vec{s}, \vec{t} \in \vec{U}$,

$$|\vec{s}| + |\vec{t}| \geq |\vec{s} \vee \vec{t}| + |\vec{s} \wedge \vec{t}|.$$

A universe \vec{U} together with such a submodular order function is called a *submodular universe*. Given a submodular universe we denote as $\vec{S}_k \subseteq \vec{U}$, for $k \in \mathbb{N}$, the separation system consisting of all separations $\vec{s} \in \vec{U}$ satisfying $|\vec{s}| < k$.

We say that a separation system \vec{S} inside a universe is *structurally submodular* (some literature omits the ‘structurally’) if, for all $\vec{s}, \vec{t} \in \vec{S}$, at least one of $\vec{s} \vee \vec{t}$ and $\vec{s} \wedge \vec{t}$ also lies in \vec{S} . Note that, if \vec{U} is a submodular universe, then every $\vec{S}_k \subseteq \vec{U}$ is structurally submodular.

2.2. The tree-of-tangles theorem (cf. [5])

A separation s is said to *distinguish* two orientations O_1 and O_2 of two, possibly distinct, separation systems inside U , if s has an orientation \vec{s} such that $\vec{s} \in O_1$ and $\vec{\bar{s}} \in O_2$. If U comes with an order function we say that such an s distinguishes O_1 and O_2 *efficiently* if there is no r with $|r| < |s|$ which distinguishes O_1 and O_2 .

A consistent orientation O of a separation system $S \subseteq U$ inside some universe U is said to be a *profile* if it satisfies the profile property:

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

For a universe U with an order function, a k -profile in \vec{U} is a profile of $S_k \subseteq U$. We say that P is a profile in \vec{U} if P is a k -profile in \vec{U} for some k . If P is a k -profile in \vec{U} , then k is the order of P .

Such a profile P is *robust* if moreover:

$$\forall \vec{s} \in P, \vec{t} \in \vec{U}: \text{ if } |\vec{s} \vee \vec{t}| < |\vec{s}| \text{ and } |\vec{s} \vee \vec{t}| < |\vec{t}|, \text{ then either } \vec{s} \vee \vec{t} \in P \\ \text{ or } \vec{s} \vee \vec{t} \in P.$$

The tree-of-tangles theorem for k -profiles states the following:

Theorem 2.2 ([5, Corollary 3.7], modified). *Let $(\vec{U}, \leq, *, \vee, \wedge, | \cdot |)$ be a submodular universe of separations. For every set \mathcal{P} of pairwise distinguishable robust regular profiles in \vec{U} there is a regular tree set $T = T(\mathcal{P}) \subseteq \vec{U}$ of separations such that:*

1. every two profiles in \mathcal{P} are efficiently distinguished by some separation in T ;
2. every separation in T efficiently distinguishes a pair of profiles in \mathcal{P} .

Note that the original statement [5, Corollary 3.7] included a third property which guaranteed that the resulting set T is invariant under automorphisms. Our methods in this paper will not allow us to guarantee this, that is why we exclude this property from our version of [5, Corollary 3.7]. For further discussion of this property, *canonicity*, see [5, 9].

Similarly, we have the tree-of-tangles theorem already mentioned in the introduction for structurally submodular separation systems which do not necessarily come in the form of an $S_k \subseteq U$:

Theorem 1.1 ([4, Theorem 6]). *Let \vec{S} be a structurally submodular separation system and \mathcal{P} a set of profiles of S . Then \vec{S} contains a tree set that distinguishes \mathcal{P} .*

2.3. Tangle-tree duality (cf. [7])

Given some set \mathcal{F} of subsets of S , an \mathcal{F} -tangle of S is a consistent orientation of S which includes no subset in \mathcal{F} . Given a submodular universe \vec{U} , we say that τ is an \mathcal{F} -tangle in \vec{U} if τ is an \mathcal{F} -tangle of some S_k . Observe that profiles are \mathcal{P} -tangles for the set \mathcal{P} of all ‘profile triples’ $\{\vec{r}, \vec{s}, (\vec{r} \vee \vec{s})^*\} \subseteq \vec{S}$.

Often we will consider sets \mathcal{F} of stars: A star in \vec{S} is a set $\sigma \subseteq \vec{S}$ such that $\vec{s} \leq \vec{t}$ for all $\vec{s}, \vec{t} \in \sigma$.

We say that a set \mathcal{F} forces a separation $\vec{s} \in \vec{S}$ if $\{\vec{s}\} \in \mathcal{F}$.

\mathcal{F} is *standard* for \vec{S} if it forces all trivial separations, that is \mathcal{F} contains all singletons $\{\vec{s}\}$ for cotrivial $\vec{s} \in \vec{S}$.

Given a tree T we denote as $\vec{E}(T)$ the set of orientations of edges of T . This set is equipped with a natural partial order where $\vec{e} \leq \vec{f}$ if and only if the unique path in T from the tail of \vec{e} to the head of \vec{f} contains both the head of \vec{e} and the tail of \vec{f} . This partial order, together with $*$ the reversal of directed edges, turns $\vec{E}(T)$ into a separation system.

Given a separation system \vec{S} , an S -tree (T, α) is a tree T together with a function $\alpha: \vec{E}(T) \rightarrow \vec{S}$ which commutes with $*$, i.e., $\alpha(\vec{e}) = \alpha(\vec{e})^*$. The S -tree is *order-respecting* if α preserves the partial order from $\vec{E}(T)$, i.e., $\alpha(\vec{e}) \leq \alpha(\vec{f})$ whenever $\vec{e} \leq \vec{f}$. For $t \in V(T)$ we denote as $\alpha(t)$ the set $\{\alpha(st) \mid s \in N(t)\}$. Given some set \mathcal{F} of subsets of S , an S -tree (T, α) is *over* \mathcal{F} if $\alpha(t) \in \mathcal{F}$ for all $t \in V(T)$.

An S -tree (T, α) is *irredundant*, if for any node $t \in V(T)$ and distinct neighbours $t', t'' \in N(t)$ we have that $\alpha(t', t) \neq \alpha(t'', t)$.

Note that, if \mathcal{F} is a set of stars, then any irredundant S -tree over \mathcal{F} is order-respecting.

Given a separation system \vec{S} inside a universe \vec{U} and $\vec{r}, \vec{s}_0 \in \vec{S}$ with $\vec{s}_0 \geq \vec{r}$ and where \vec{r} is non-degenerate and non-trivial in \vec{S} , the *shifting map* $f \downarrow_{\vec{s}_0}^{\vec{r}}$ is defined by letting, for every $\vec{s} \geq \vec{r}$,

$$f \downarrow_{\vec{s}_0}^{\vec{r}}(\vec{s}) = \vec{s} \vee \vec{s}_0 \quad \text{and} \quad f \downarrow_{\vec{s}_0}^{\vec{r}}(\vec{s}) = (\vec{s} \vee \vec{s}_0)^*.$$

This map is defined on $\vec{S}_{\geq \vec{r}} \setminus \{\vec{r}\}$, where $S_{\geq \vec{r}}$ is the set of all separations $t \in S$ which have an orientation \vec{t} with $\vec{t} \geq \vec{r}$, and $\vec{S}_{\geq \vec{r}}$ is the set of all orientations of separations in $S_{\geq \vec{r}}$.

For an irredundant S -tree (T, α) over some set of stars with $\{\vec{r}\} = \alpha(x)$, for some leaf x of T , we write

$$\alpha_{x, \vec{s}_0} := f \downarrow_{\vec{s}_0}^{\vec{r}} \circ \alpha.$$

The resulting new tree $(T, \alpha_{x, \vec{s}_0})$ is called the *shift of (T, α) from \vec{r} to \vec{s}_0* if the leaf x is the only one which has $\alpha(x) = \{\vec{r}\}$.

Given a separation system \vec{S} inside a universe \vec{U} and a star $\sigma \subseteq \vec{S}$ a *shift of σ (to some $\vec{s}_0 \in \vec{S}$)* is a star of the form

$$\sigma_{\vec{x}}^{\vec{s}_0} := \{\vec{x} \vee \vec{s}_0\} \cup \{\vec{y} \wedge \vec{s}_0 \mid \vec{y} \in \sigma \setminus \{\vec{x}\}\},$$

where $\vec{x} \in \sigma$. Note that if, for some $\vec{r} \in \vec{S}$, we have $\vec{x} \geq \vec{r}$, then $\sigma_{\vec{x}}^{\vec{s}_0}$ is the image of σ under $f \downarrow_{\vec{s}_0}^{\vec{r}}$.

A separation \vec{s} *emulates* \vec{r} in \vec{S} if $\vec{s} \geq \vec{r}$ and for every $\vec{t} \in \vec{S} \setminus \{\vec{r}\}$ with $\vec{t} \geq \vec{r}$ we have $\vec{s} \vee \vec{r} \in \vec{S}$. The separation \vec{s} emulates \vec{t} in \vec{S} for \mathcal{F} if additionally for every star $\sigma \in \mathcal{F}$ with $\vec{r} \notin \sigma$ and every $\vec{x} \in \sigma$ with $\vec{x} \geq \vec{r}$ we have $\sigma_{\vec{x}}^{\vec{s}} \in \mathcal{F}$.

Note that for an irredundant S -tree (T, α) over some set of stars \mathcal{F} with $\{\vec{r}\} = \alpha(x)$, for some leaf x of T , the shift from \vec{r} to \vec{s}_0 is again an S -tree over \mathcal{F} if \vec{s}_0 emulates \vec{r} in \vec{S} for \mathcal{F} .

A separations system S is *separable* if for any two nontrivial non-degenerate separations $\vec{r}_1, \vec{r}_2 \in \vec{S}$ with $\vec{r}_1 \leq \vec{r}_2$ there exists a separation $\vec{s}_0 \in \vec{S}$, with $\vec{r}_1 \leq \vec{s}_0 \leq \vec{r}_2$ such that \vec{s}_0 emulates \vec{r}_1 in S and \vec{s}_0 emulates \vec{r}_2 in S . The separation system S is *\mathcal{F} -separable* if we can choose, for any two such \vec{r}_1 and \vec{r}_2 which are nontrivial non-degenerate and not forced by \mathcal{F} , such an \vec{s}_0 so that \vec{s}_0 emulates \vec{r}_1 in S for \mathcal{F} and \vec{s}_0 emulates \vec{r}_2 in S for \mathcal{F} .

The abstract tangle-tree duality theorem now states the following:

Theorem 1.2 (Tangle-tree duality theorem [7, Theorem 4.3]). *Let U be a universe of separations containing a finite separation system $S \subseteq U$ and let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars such that \mathcal{F} is standard for \vec{S} and \vec{S} is \mathcal{F} -separable. Then exactly one of the following statements holds:*

- *there is an \mathcal{F} -tangle of S ;*
- *there is an S -tree over \mathcal{F} .*

If, in the following, we speak of *the* duality theorem, we mean Theorem 1.2.

The condition of \mathcal{F} -separability is sometimes split into two parts which, in sum, are stronger: Firstly, that S is separable and secondly that \mathcal{F} is *closed under shifting*, that is, every shift σ' of a star $\sigma \in \mathcal{F}$ is also in \mathcal{F} if $\sigma' \subseteq \vec{S}$. (Compare [4, Lemma 12].)

We shall need the following additional lemmas from the literature:

Lemma 2.3 ([7, Lemma 2.1]). *Every irredundant S -tree (T, α) over stars is order-respecting. In particular, $\alpha(\vec{E}(T))$ is a nested set of separations in \vec{S} .*

Lemma 2.4 ([7, Lemma 2.2]). *Let (T, α) be an irredundant S -tree over a set \mathcal{F} of stars. Let e, f be distinct edges of T with orientations $\vec{e} < \vec{f}$ such that $\alpha(\vec{e}) = \alpha(\vec{f}) =: \vec{r}$. Then \vec{r} is trivial.*

In particular, T cannot have distinct leaves associated with the same star $\{\vec{r}\}$ unless \vec{r} is trivial.

Lemma 2.5 ([7, Lemma 2.3]). *If (T, α) is an S -tree over \mathcal{F} , possibly redundant, then T has a subtree T' such that (T', α') is an irredundant S -tree over \mathcal{F} , where α' is the restriction of α to $\vec{E}(T')$. If (T, α) is rooted at a leaf*

x and T has an edge, then T' can be chosen so as to contain x and e_x , the edge incident to x in T .

Lemma 2.6 ([7, Lemma 2.4]). *Let (T, α) be an S -tree over a set \mathcal{F} of stars, rooted at a leaf x . Assume that T has an edge, and that $\vec{r} = \alpha(\vec{e}_x)$ is nontrivial. Then T has a minor T' containing x and e_x such that (T', α') , where $\alpha' = \alpha \upharpoonright \vec{E}(T')$, is a tight and irredundant S -tree over \mathcal{F} .*

For every such (T', α') the edge \vec{e}_x is the only edge $\vec{e} \in \vec{E}(T')$ with $\alpha(\vec{e}) = \vec{r}$.

Lemma 2.7 ([4, Lemma 13]). *Let \vec{U} be a universe of separations and $\vec{S} \subseteq \vec{U}$ a structurally submodular separation system. Then \vec{S} is separable.*

Moreover, we shall need a variant of [7, Lemma 4.2] which follows with the exact same proof:

Lemma 2.8 ([7]). *Let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars. Let (T, α) be a tight and irredundant S -tree with at least one edge, over some set of stars, and rooted at a leaf x . Assume that $\vec{r} := \alpha(\vec{e}_x)$ is nontrivial and non-degenerate, let $\vec{s}_0 \in \vec{S}$ emulate \vec{r} in \vec{S} for \mathcal{F} , and consider $\alpha' := \alpha_{x, \vec{s}_0}$. Then (T, α') is an order-respecting S -tree in which $\{\vec{s}_0\}$ is a star associated with x but with no other leaf of T . Moreover $\alpha'(t) \in \mathcal{F}$ for all $t \neq x$ with $\alpha(t) \in \mathcal{F}$.*

The only difference in the statement between Lemma 2.8 and [7, Lemma 4.2] is that [7, Lemma 4.2] requires that (T, α) is an S -tree over \mathcal{F} , whereas we only require (T, α) to be an S -tree over some set of stars. Consequently, in [7, Lemma 4.2] it is shown that then (T, α') is an S -tree over $\mathcal{F} \cup \{\{\vec{s}_0\}\}$ whereas we only conclude that $\alpha'(t) \in \mathcal{F}$ whenever $\alpha(t) \in \mathcal{F}$.

2.4. Splices in submodular universes

In addition to the existing terminology, we shall need the following new concept, which has already been considered in [6], but has not been given a name there: In a submodular universe \vec{U} a separation \vec{s} is a *splice* for a separation \vec{r} with $\vec{r} \leq \vec{s}$ if there is no separation \vec{t} with $\vec{r} \leq \vec{t} \leq \vec{s}$ and $|t| < |s|$. A *splice between* two separations \vec{r} and \vec{s} with $\vec{r} \leq \vec{s}$ is one of minimum order among all \vec{t} with $\vec{r} \leq \vec{t} \leq \vec{s}$.

These splices are good choices for proving separability due to the next lemma. It follows directly from the proof of Lemma 3.4 of [6] which, phrased in our terminology, considers a splice between two separations. We recapitulate the main argument of this proof below.

Lemma 2.9 ([6]). *Consider $\vec{S}_k \subseteq \vec{U}$ in a submodular universe. If $\vec{s} \in \vec{S}_k$ is a splice for $\vec{r} \in \vec{S}_k$, then, for every $\vec{t} \in \vec{U}$ with $\vec{t} \geq \vec{r}$, the order of $\vec{t} \vee \vec{s}$ is at most the order of \vec{t} . In particular, \vec{s} emulates \vec{r} in \vec{S}_k .*

Proof sketch, cf. [6, Lemma 3.4]. If the order of $\vec{t} \vee \vec{s}$ were greater than the order of \vec{t} , then, by submodularity, the order of $\vec{t} \wedge \vec{s}$ would be less than the order of \vec{s} . However, by the fish Lemma 2.1, $\vec{r} \leq \vec{t} \wedge \vec{s} \leq \vec{s}$ and this contradicts the fact that \vec{s} is a splice for \vec{r} . \square

This lemma then directly implies the ultimate statement of [6, Lemma 3.4]:

Lemma 2.10 ([6, Lemma 3.4]). *Every $\vec{S}_k \subseteq \vec{U}$ in a submodular universe is separable.*

3. Structurally submodular separation systems

In this section we will prove the first tree-of-tangles theorem of this paper. It is a theorem for regular profiles, all of the same structurally submodular separation system, and states as follows:

Theorem 3.1. *Let \vec{S} be a structurally submodular separation system. Then S contains a nested set that distinguishes the set of regular profiles of S .*

By itself Theorem 3.1 is nothing special; indeed, it is a slight weakening of Theorem 1.1, which asserts the same but without requiring the profiles to be regular. In this case the ingredients of the proof are more interesting than its result: we shall obtain Theorem 3.1 as a direct consequence of Theorem 1.2.

So let \vec{S} be a structurally submodular separation system inside some universe \vec{U} . Since we are interested in the regular profiles of S we may assume that S has no degenerate elements. Our strategy will be as follows: we shall construct a set $\mathcal{F} \subseteq 2^{\vec{U}}$ for which there is no \mathcal{F} -tangle of S , and so that every element of \mathcal{F} is included in at most one regular profile of S . If we can achieve this, then Theorem 1.2 applied to this set \mathcal{F} will yield an S -tree over \mathcal{F} . The set N of edge labels of this S -tree (T, α) will then be the desired nested set distinguishing all regular profiles of S : each regular profile P of S orients the edges of T and hence includes a star σ of the form $\alpha(t)$ for some $t \in V(T)$. By choice of \mathcal{F} this σ is included in no other regular profile of S , which means that it distinguishes P from all other profiles.

To construct this set \mathcal{F} , first let \mathcal{P} be the set of all ‘profile triples’ in \vec{S} : the set of all $\{\vec{r}, \vec{s}, (\vec{r} \vee \vec{s})^*\} \subseteq \vec{S}$. For a consistent orientation of S it is then

equivalent to be a profile of S and to be a \mathcal{P} -tangle. Furthermore let \mathcal{C} be the set of all $\{\vec{s}\}$ with $\vec{s} \in \vec{S}$ co-small. Finally, let \mathcal{M} consist of each of the sets $\max P$ of maximal elements of P for each regular profile P of S . We then take

$$\mathcal{F} := \mathcal{P} \cup \mathcal{C} \cup \mathcal{M}.$$

With these definitions the regular profiles of S are precisely its $(\mathcal{P} \cup \mathcal{C})$ -tangles; and there are no \mathcal{F} -tangles of S since each regular profile P of S includes $\max P \in \mathcal{M} \subseteq \mathcal{F}$. If this \mathcal{F} were a set of stars and if we could feed this \mathcal{F} to Theorem 1.2, we would receive an S -tree over \mathcal{F} and the edge labels of this S -tree would be our desired nested set, since each element of \mathcal{F} is included in at most one regular profile of S : indeed, the regular profiles of S have no subsets in \mathcal{P} or \mathcal{C} , and each element $\max P \in \mathcal{M}$ is included only in P itself.

Unfortunately, we are still some way off from plugging \mathcal{F} into Theorem 1.2: we need to ensure that \mathcal{F} is a set of stars that is standard for S and that S is \mathcal{F} -separable. Out of these the second and one half of the third are easy: \mathcal{F} is standard for S since $\mathcal{C} \subseteq \mathcal{F}$ is, and S is separable by Lemma 2.7.

We thus need to show that S is not only separable but \mathcal{F} -separable. Unfortunately our current set \mathcal{F} is not even a set of stars yet. However, in [3] a solution was laid out for this exact situation: a series of lemmas from [3] shows that we can simply *make* \mathcal{F} a set of stars and close it under shifting without altering the set of \mathcal{F} -tangles of S .

The way to do this is as follows. Given two elements \vec{r} and \vec{s} of some set $\sigma \subseteq \vec{S}$, by submodularity, either $\vec{r} \wedge \vec{s}$ or $\vec{r} \vee \vec{s}$ must lie in \vec{S} . *Uncrossing* \vec{r} and \vec{s} in σ then means to replace either \vec{r} by $\vec{r} \wedge \vec{s}$ or \vec{s} by $\vec{r} \vee \vec{s}$, depending on which of these two lies in \vec{S} . (Structural submodularity ensures that at least one of them does.) Uncrossing all pairs of elements of σ in turn yields a star σ^* , which we call an *uncrossing* of σ . (Note that σ^* is not in general unique since it depends on the order in which one uncrosses the elements of σ .) It is then easy to see that a regular profile of S includes σ if and only if it includes σ^* :

Lemma 3.2 ([3, Lemma 11]). *If a regular profile of S includes an uncrossing of some set, it also includes that set.*

Conversely, if a regular consistent orientation of S includes some set, it also includes each uncrossing of that set.

Let us write \mathcal{F}^* for the set of all uncrossings of elements of \mathcal{F} . Then \mathcal{F}^* is a set of stars that is standard for S . We are still not done, however, since \mathcal{F}^* need not be closed under shifting. We can fix this in a similar manner though.

Just as for uncrossings it is not hard to show that the inclusion of a star's shift in a regular profile implies that star's inclusion:

Lemma 3.3 ([3, Lemma 13]). *If a regular profile of S includes a shift of some star, it also includes that star.*

In [3] the definition of a shift of a star contains additional technical assumptions on σ and $\vec{\sigma}_0$, keeping in line with the precise assumptions of Theorem 1.2. However the proof of Lemma 3.3 does not necessitate this, and neither does its application.

Lemma 3.3 says that if we close \mathcal{F}^* under shifting we, again, do not alter the set of \mathcal{F}^* -tangles of S . Formally, set $\mathcal{G}_0 = \mathcal{F}^*$, and for $i \geq 1$ let \mathcal{G}_i be the set of all shifts of stars in \mathcal{G}_{i-1} . Write $\hat{\mathcal{F}}^* := \bigcup_{i \in \mathbb{N}} \mathcal{G}_i$. Then by Lemma 3.3 the $\hat{\mathcal{F}}^*$ -tangles of S are precisely its \mathcal{F}^* -tangles, which is to say that there are no $\hat{\mathcal{F}}^*$ -tangles of S . Moreover this set $\hat{\mathcal{F}}^*$ still has the property that each star in it is included in at most one regular profile: let us say that $\hat{\sigma}^* \in \hat{\mathcal{F}}^*$ originates from $\sigma \in \mathcal{F}$ if $\hat{\sigma}^*$ can be obtained by a series of shifts from an uncrossing of σ . Lemma 3.2 and 3.3 then say that if $\hat{\sigma}^* \subseteq P$ for a regular profile P , and $\hat{\sigma}^*$ originates from $\sigma \in \mathcal{F}$, then $\sigma \subseteq P$. Since the only element of \mathcal{F} which P includes is $\max P$, this implies that no other regular profile of S includes $\hat{\sigma}^*$.

We can thus formally prove Theorem 3.1:

Proof of Theorem 3.1. Define \mathcal{P} , \mathcal{C} , \mathcal{M} , \mathcal{F} , \mathcal{F}^* , and $\hat{\mathcal{F}}^*$ as above. Then $\hat{\mathcal{F}}^*$ is standard for S since $\mathcal{C} \subseteq \hat{\mathcal{F}}^*$, and closed under shifting by construction. By Lemma 2.7 S is separable. Together this gives that S is \mathcal{F} -separable. Hence we can apply the tangle-tree duality theorem 1.2 to obtain either an $\hat{\mathcal{F}}^*$ -tangle of S or an S -tree over $\hat{\mathcal{F}}^*$.

We claim that the first is impossible. For suppose that P is some $\hat{\mathcal{F}}^*$ -tangle of S . From $\mathcal{C} \subseteq \hat{\mathcal{F}}^*$ we know that P is a regular and consistent orientation of S . If P has the profile property (P), then we could derive a contradiction from Lemma 3.2 and 3.3 since S has no \mathcal{F} -tangle. On the other hand, if P is not a profile, then P includes some set $\sigma \in \mathcal{P}$. By the second part of Lemma 3.2 P then also includes some (in fact: each) uncrossing of σ and hence a set in $\mathcal{F}^* \subseteq \hat{\mathcal{F}}^*$, contrary to its status as an $\hat{\mathcal{F}}^*$ -tangle.

So let (T, α) be the S -tree over $\hat{\mathcal{F}}^*$ returned by Theorem 1.2, which we may assume to be irredundant (Lemma 2.5). Let \vec{N} be the image of α . Then N is a nested subset of S (Lemma 2.3). Let us show that N distinguishes all regular profiles of S . Since (T, α) is an S -tree over $\hat{\mathcal{F}}^*$ each consistent orientation of S includes some star $\hat{\sigma}^* \in \hat{\mathcal{F}}^* \cap 2^{\vec{N}}$. In particular if P is a regular profile of S , then P includes some $\hat{\sigma}^* \in \hat{\mathcal{F}}^* \cap 2^{\vec{N}}$. Since

the only element of \mathcal{F} which P includes is $\max P$, this $\hat{\sigma}^*$ must originate from $\max P$. Consequently no other regular profile of S can include $\hat{\sigma}^*$, since none of them include $\max P$. Thus $\hat{\sigma}^*$ distinguishes P from every other regular profile of S . Since P was arbitrary this shows that N distinguishes all regular profiles of S . \square

Let us make some remarks on this proof of Theorem 3.1. First, in the definition of \mathcal{F} , we could have used other sets \mathcal{M} : the only properties of \mathcal{M} that we used is that every regular profile of S contains some set from \mathcal{M} , and that no element of \mathcal{M} is included in more than one such regular profile. We will put this observation to good use in Section 4, where we will make a more refined choice for \mathcal{M} than simply collecting the sets of maximal elements from each profile.

Second, with the approach shown here it is not easy to strengthen Theorem 3.1 to the level of Theorem 1.1 by dropping the assumption of regularity, since Lemma 3.3 cannot do without this regularity.

In the remainder of this section we will show a more direct version of the proof presented above. This proof will be the guiding principle by which we will approach the issues of efficiency and profiles of differing order in Sections 5 and 7.

The core idea is that one can take as \mathcal{F} the set of all stars that are included in at most one regular profile of S . An S -tree over this set \mathcal{F} would immediately lead to the desired nested set distinguishing all regular profiles. Moreover this \mathcal{F} is standard for S since $\mathcal{C} \subseteq \mathcal{F}$. To obtain this S -tree over \mathcal{F} from Theorem 1.2 one would only need to show two things, namely that \vec{S} is \mathcal{F} -separable and that there is no \mathcal{F} -tangle of S . The first of these amounts to Lemma 3.3; the second requires the two insights that every \mathcal{F} -avoiding consistent orientation is a regular profile, and that each regular profile of S includes some star in \mathcal{F} , both of which retrace some steps of Lemma 3.2.

Lemma 3.4. *Let $\vec{S} \subseteq \vec{U}$ be a structurally submodular separation system and let P be a profile of \vec{S} . There exists a star $\sigma \subseteq P$ such that no other profile of \vec{S} includes σ .*

Proof. Let $\sigma \subseteq P$ be a star which minimizes the number of profiles which include σ . Suppose for a contradiction that there exists a profile $P' \neq P$ with $\sigma \subseteq P'$. Some separation s , say, distinguishes P from P' . Clearly s crosses some element of σ .

Suppose that, subject to the above, σ and s are chosen so that the number of separations in σ that s crosses is minimum. Let $\vec{t} \in \sigma$ be a separation that s crosses. If either of the corner separations $\vec{t} \vee \vec{s}$ or $\vec{t} \vee \vec{\bar{s}}$

was in \vec{S} , then, by the profile property, it would distinguish P and P' . It would also, by the fish Lemma 2.1, cross one less separation in σ than s does, contradicting the choice of s .

So by submodularity the corner separations $\vec{t} \wedge \vec{s}$ and $\vec{t} \wedge \vec{s}$ are in \vec{S} . Note that, by the profile property, any profile including

$$\sigma' := \sigma \setminus \{\vec{t}\} \cup \{\vec{t} \wedge \vec{s}, \vec{t} \wedge \vec{s}\}$$

also includes σ . Consequently σ' together with s are a better choice than σ and s , a contradiction. \square

Lemma 3.5. *Given any set \mathcal{P} of profiles of \vec{S} , every consistent orientation O of \vec{S} which is not a profile in \mathcal{P} contains a star σ which is not contained in any profile in \mathcal{P} .*

Proof. Since O is not a profile in \mathcal{P} there is, for every profile P in \mathcal{P} , a separation s such that $\vec{s} \in O$ but $\vec{s} \in P$. Pick a set $N \subseteq O$ which contains one such separation for every profile in \mathcal{P} and is, subject to this, \leq -minimal: That is, there is no other such set N' together with an injective function $\alpha : N' \rightarrow N$ satisfying $\vec{s} \leq \alpha(\vec{s}')$ for all $\vec{s}' \in N'$.

If N is a nested set, then N contains the desired star, so suppose that $\vec{s}, \vec{t} \in N$ cross. By submodularity we may suppose, after possibly renaming \vec{s} and \vec{t} , that $\vec{s} \wedge \vec{t} \in S$ and thus, by consistency, $\vec{s} \wedge \vec{t} \in O$. We claim that $N \setminus \{\vec{s}\} \cup \{\vec{s} \wedge \vec{t}\}$ is also a candidate for N , contradicting the \leq -minimality. So suppose that $N \setminus \{\vec{s}\} \cup \{\vec{s} \wedge \vec{t}\}$ does not contain a separation \vec{r} such that $\vec{r} \in P$, say. Then clearly $\vec{s} \in P$ and $\vec{t} \in P$, thus, by the profile property $\vec{s} \vee \vec{t} \in P$ which is precisely such an \vec{r} , a contradiction. \square

We are now ready to give a proof of Theorem 3.1 without resorting to Lemma 3.2:

Theorem 3.1. *Let \vec{S} be a structurally submodular separation system. Then S contains a nested set that distinguishes the set of regular profiles of S .*

Direct Proof. Let \mathcal{P} be the set of regular profiles of S . Let $\mathcal{F}_{\mathcal{P}} \subseteq 2^{\vec{S}}$ consist of all stars $\sigma \subseteq \vec{S}$ for which one of the following is true:

- (i) No profile in \mathcal{P} includes σ , or
- (ii) Precisely one profile in \mathcal{P} includes σ .

This $\mathcal{F}_{\mathcal{P}}$ is, by Lemma 3.3, closed under shifting: any shift of a star contained in at most one profile is again contained in at most one profile. The set $\mathcal{F}_{\mathcal{P}}$ is also standard for \vec{S} , since cosmall separations are contained in no regular profile.

By Theorem 1.2 there either exists an S -tree over $\mathcal{F}_{\mathcal{P}}$, or an $\mathcal{F}_{\mathcal{P}}$ -tangle of S . In the former case we obtain the desired nested set. For the latter case observe that every $\mathcal{F}_{\mathcal{P}}$ -tangle P , say, is a regular profile: By Lemma 3.5 every consistent orientation which avoids $\mathcal{F}_{\mathcal{P}}$ is a profile and if P would not be regular, it would contain a cosmall separation \vec{s} which is impossible, since $\{\vec{s}\} \in \mathcal{F}_{\mathcal{P}}$. So by Lemma 3.4 there exists a star $\sigma \subseteq P$ which every profile other than P avoids. In particular $\sigma \in \mathcal{F}_{\mathcal{P}}$, which contradicts the fact that P is an $\mathcal{F}_{\mathcal{P}}$ -tangle. \square

4. Application: degrees in trees of tangles

In this section we are going to see that our proof of Theorem 3.1 in Section 3 has one advantage over the usual, more direct proofs of Theorem 3.1 from [4, 10]: It allows us to easily control the maximum degree of the resulting tree. More precisely: Let S be a structurally submodular separation system and P a regular profile of S . In this section we answer the following question: over all trees of tangles that distinguish all regular profiles of S , how low can the degree of the node containing P in those trees of tangles be?

Let us first make this notion of degree in a tree of tangles formal. For the purposes of this application only, a *tree of tangles (for S)* is an irredundant S -tree (T, α) whose set of edge labels distinguishes all regular profiles of S . For a regular profile P of S and a tree of tangles (T, α) , the *node of P in T* is the unique sink of the orientation of T 's edges induced by P , and the *degree of P in (T, α)* is the degree of this node.

Our question is thus: what is the minimum degree of P in (T, α) over all trees of tangles (T, α) ?

A lower bound for this degree can be established as follows. Let $\delta(P)$ denote the minimal size of a set of separations which distinguishes P from all other regular profiles of S . If t is the node of P in some tree of tangles (T, α) , then $\alpha(t)$ is such a set of separations which distinguishes P from all other regular profiles of S ; thus, the degree of P in every tree of tangles (T, α) is at least $\delta(P)$.

We show that this lower bound can be achieved: there is a tree of tangles (T, α) for S in which P has degree exactly $\delta(P)$. In fact (T, α) will be optimal in this sense not just for P , but for all regular profiles of S simultaneously. Additionally the degrees of those nodes of (T, α) that are not the node of some regular profile will not be unreasonably high: the maximum degree of T will be attained in some profiles' node.

Theorem 4.1. *Let S be a structurally separation system. Then there is a tree of tangles (T, α) for S in which each regular profile P of S has degree*

exactly $\delta(P)$. Furthermore, if $\Delta(T) > 3$, then $\Delta(T) = \delta(P)$ for some regular profile P of S .

To prove Theorem 4.1 we will follow the first proof of Theorem 3.1, making a more refined choice of \mathcal{M} , and utilise the fact that uncrossing and shifting a set cannot increase its size.

We will later see an example of a structurally submodular separation system in which $\delta(P) \leq 2$ for every profile P but $\Delta(T) = 3$ for every tree of tangles T ; this will demonstrate that the last assertion of Theorem 4.1 is optimal in that regard.

Observe further that the set of maximal elements of a profile P is a set which distinguishes P from every other profile of S . (In fact, the maximal elements of P distinguish P from every other consistent orientation of S .) Therefore $\delta(P) \leq |\max P|$ and hence the degree of P in the tree of tangles from Theorem 4.1 is at most $|\max P|$.

Let us now prove Theorem 4.1:

Proof of Theorem 4.1. For each regular profile P of S pick a subset $D_P \subseteq P$ of size $\delta(P)$ which distinguishes P from every other regular profile of S . Let \mathcal{D} be the set of these D_P . Define \mathcal{P} and \mathcal{C} as in the proof of Theorem 3.1, and set

$$\mathcal{F} := \mathcal{P} \cup \mathcal{C} \cup \mathcal{D}.$$

From here, define \mathcal{F}^* and $\hat{\mathcal{F}}^*$ just as in Theorem 3.1 and follow the same proof. The result is an S -tree over $\hat{\mathcal{F}}^*$, which we may assume to be irredundant and hence a tree of tangles for S .

Now let P be a regular profile of S , let t be the node of P in T , and $\hat{\sigma}^* := \alpha(t)$. As in the proof of Theorem 3.1 the only element of \mathcal{F} from which $\hat{\sigma}^*$ can originate is D_P . Since uncrossing and shifting D_P cannot increase its size we have $|\hat{\sigma}^*| \leq |D_P| = \delta(P)$. Conversely we have $|\hat{\sigma}^*| \geq \delta(P)$ since $\hat{\sigma}^*$ distinguishes P from all other regular profiles. Thus the degree of P in (T, α) is indeed $\delta(P)$.

Finally, if $\Delta(T) > 3$, the maximum degree of T is attained in some node t whose associated star $\alpha(t)$ originates from some $D_P \in \mathcal{D}$, since all elements of $\hat{\mathcal{F}}^*$ originating from elements of \mathcal{P} or \mathcal{C} have size at most three. As above we thus have $|\alpha(t)| \leq |D_P| = \delta(P)$, giving $\Delta(T) = \delta(P)$. \square

Let us see an example showing that we cannot guarantee to find T with maximum degree less than three even if all regular profiles of S have $\delta(P) \leq 2$:

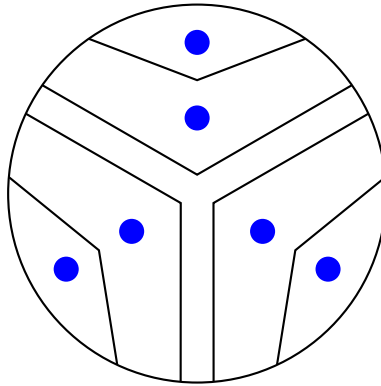


Figure 1: A ground-set and system of bipartitions.

Example 4.2. Let V consist of the six points in Fig. 1, and S be the separation system given by the six outlined bipartitions of V together with $\{\emptyset, V\}$. (That is, \vec{S} contains (A, B) and (B, A) for each of these bipartitions $\{A, B\}$. We have $(A, B) \leq (C, D) :\Leftrightarrow A \subseteq C$, and $(A, B)^* = (B, A)$. Compare [6].) The regular profiles of S correspond precisely to the six elements of V : each $v \in V$ induces a profile of S by orienting each bipartition towards v , and conversely each profile of S is of this form. Each profile P has at most two maximal elements, giving $\delta(P) \leq 2$. However, every tree of tangles for S must contain the outer three bipartitions and hence have a maximum degree of at least three.

5. Efficient distinguishers

Often our structurally submodular separation system \vec{S} is actually an \vec{S}_k , the set of all separations of order less than k , of some submodular universe \vec{U} . In this case we are not just interested in a nested set of separations which distinguishes all profiles, but one which does so *efficiently*, that is, for any two profiles it contains a distinguishing separation of minimum possible order. In this section we are going to see how this can be achieved for regular profiles of a fixed \vec{S}_k utilising the duality theorem together with a separate application of its core mechanism: shifting S -trees.

We will prove this theorem:

Theorem 5.1. *Let \vec{U} be a submodular universe and let \mathcal{P} be a set of regular profiles of \vec{S}_k . Then there exists a nested set $N \subseteq S_k$ efficiently distinguishing all the profiles in \mathcal{P} .*

Our approach is similar to the one of the direct proof in Section 3, but we shall restrict our set of stars so that they do not interfere with efficiency.

Consider a nested set of separations which distinguishes all profiles efficiently and, subject to this, is \subseteq -minimal. Every profile P induces an orientation of this set, and the maximal elements of this orientation form a star. The separations in this star are, in a way, ‘well connected’ to the profile. We make this a condition on the stars we consider. For a star σ and a profile P , we say that σ has the property $\text{Eff}(P)$ if the following holds:

$$(\text{Eff}(P)) \quad \nexists \vec{s} \in \sigma \text{ and } \vec{s}' \in P: \vec{s} \leq \vec{s}' \text{ and } |s'| < |s|.$$

This condition ensures that, for two profiles P and P' , a star σ with property $\text{Eff}(P)$ containing \vec{s} , and a star σ' with property $\text{Eff}(P')$ containing \vec{s} , the separation \vec{s} needs to be an efficient P - P' -distinguisher. For if \vec{s} is not efficient, consider an efficient P - P' -distinguisher $\vec{r} \in P$. Then \vec{r} cannot be nested with \vec{s} , since $\vec{s} \leq \vec{r}$ would contradict property $\text{Eff}(P)$ whereas $\vec{r} \leq \vec{s}$ would contradict property $\text{Eff}(P')$. But \vec{r} cannot cross \vec{s} either: if it did, we would have either $|\vec{r} \vee \vec{s}| < |\vec{s}|$ or $|\vec{r} \wedge \vec{s}| < |\vec{s}|$ by submodularity, again contradicting property $\text{Eff}(P)$ or $\text{Eff}(P')$, respectively.

Property $\text{Eff}(P)$ is preserved under taking shifts:

Lemma 5.2. *Let $\vec{s} \in \vec{S}_k$ be a splice for $\vec{r} \in \vec{S}_k$ and let $\sigma \subseteq \vec{S}_k$ be a star with some $\vec{x} \in \sigma$ with $\vec{x} \geq \vec{r}$. If a profile P contains both σ and $\sigma' := \sigma_{\vec{x}}^{\vec{s}}$ and σ has property $\text{Eff}(P)$, then also σ' has property $\text{Eff}(P)$.*

Proof. Suppose for a contradiction that σ' does not have property $\text{Eff}(P)$, that is, above some $\vec{t} \wedge \vec{s} \in \sigma'$, where $\vec{t} \in \sigma$, there is a separation $\vec{t}' \in P$ of lower order than $\vec{t} \wedge \vec{s}$.

We will first show that we may assume $\vec{t}' \leq \vec{t}$. Since \vec{s} is a splice for \vec{r} we have $|\vec{s} \wedge \vec{t}| \geq |\vec{s}|$, and thus by submodularity $|\vec{s} \wedge \vec{t}'| \leq |\vec{t}'|$. So if $\vec{t}' > \vec{t}$, then this contradicts the assertion that σ has property $\text{Eff}(P)$. If however \vec{t}' crosses \vec{t} , then, by the profile property of P and property $\text{Eff}(P)$ of σ , the supremum $\vec{t}' \vee \vec{t}$ has at least the order of \vec{t} . By submodularity then $\vec{t}' \wedge \vec{t}$ has at most the order of \vec{t}' . This is also a separation in P which is above $\vec{t} \wedge \vec{s}$ and of lower order than $\vec{t} \wedge \vec{s}$, so we may consider it instead.

Now, since \vec{s} is a splice for \vec{r} we have that $|\vec{t}' \wedge \vec{s}| \geq |\vec{s}|$, so by submodularity $\vec{t}' \wedge \vec{s}$ has at most the order of \vec{t}' . But this $\vec{t}' \wedge \vec{s}$ is the same as $\vec{t} \wedge \vec{s}$ since $\vec{t} \geq \vec{t}' \geq \vec{t} \wedge \vec{s}$. So we have $|\vec{t} \wedge \vec{s}| \leq |\vec{t}'|$, which contradicts the assumption that $|\vec{t}'| < |\vec{t} \wedge \vec{s}|$. \square

We define \mathcal{F}_e as the set of all stars $\sigma \subseteq \vec{S}_k$ which are contained in at most one profile in \mathcal{P} and which, if they are contained in a profile $P \in \mathcal{P}$, fulfill property $\text{Eff}(P)$.

From Lemma 5.2 and 3.3 immediately we obtain the following corollary:

Corollary 5.3. S_k is \mathcal{F}_e -separable.

However, an S -tree over \mathcal{F}_e does not necessarily give rise to an efficient distinguisher set for \mathcal{P} because we make no assumptions on those stars which are not contained in any profile. Our proof of Theorem 5.1 will need to make additional arguments on why an efficient such tree exists.

It would be much more elegant if we could introduce a condition, similar to $\text{Eff}(\cdot)$, on the stars which are in no profile, so as to guarantee that any S_k -tree over these stars is as desired. However, all possible such properties that the authors could come up with failed to give \mathcal{F} -separability and there is reason to believe that such a solution is not possible: The critical part in the proof of Theorem 5.1 will make a global argument, specifically that of two shifts of one separation one is an efficient distinguisher. Separability on the other hand is defined in terms of each individual shift of a star.

For this section’s analogue of Lemma 3.4, we define the *fatness* of a star σ as the tuple $(n_{k-1}, n_{k-2}, \dots, n_1, n_0)$, where n_i is the number of separations of order i in σ . We will consider the lexicographic order on the fatness of stars.

Lemma 5.4. *Given a set \mathcal{P} of regular profiles of \vec{S}_k , every profile $P \in \mathcal{P}$ includes a star in \mathcal{F}_e .*

Proof. By Lemma 3.4 P includes a star which is contained only in P . Take such a star σ which has lexicographically minimal fatness and suppose for a contradiction that σ does not have property $\text{Eff}(P)$. So take $\vec{s} \in \sigma$ and $\vec{r} \in P$ with $\vec{s} \leq \vec{r}$ and $|r| < |s|$. Among the possible choices for \vec{r} , let \vec{r} be one which crosses as few separations in σ as possible. If \vec{r} were nested with σ , then the maximal elements of $\sigma \cup \{\vec{r}\}$ would form a star of lower fatness, thus we may suppose that \vec{r} crosses some $\vec{x} \in \sigma$.

By the choice of \vec{r} , the corner separations $\vec{r} \vee \vec{x}$ and $\vec{r} \wedge \vec{x}$ must have strictly higher order than $|r|$ since both are $\geq \vec{s}$. Thus, by submodularity, the corner separations $\vec{r} \wedge \vec{x}$ and $\vec{r} \vee \vec{x}$ have strictly lower order than $|x|$. Now the star $\sigma' := \sigma \setminus \{\vec{x}\} \cup \{\vec{r} \wedge \vec{x}, \vec{r} \vee \vec{x}\}$ has a lower fatness. This star is still contained in P by consistency and in no other profile, since every profile which includes σ' also includes σ by the profile property applied with \vec{x} and \vec{r} . This contradicts the choice of σ . \square

We are now able to prove Theorem 5.1:

Proof of Theorem 5.1. We may apply Theorem 1.2 for \mathcal{F}_e since \vec{S}_k is \mathcal{F}_e -separable by Corollary 5.3 and \mathcal{F}_e is standard since cotrivial separations are

not contained in any regular profile. From this theorem we cannot get an \mathcal{F}_e -tangle: such a tangle cannot be a profile in \mathcal{P} by Lemma 5.4, and Lemma 3.5 states that every consistent orientation which is not a profile in \mathcal{P} includes a star which is not contained in any profile in \mathcal{P} , but each of these stars is contained in \mathcal{F}_e , so no such orientation is an \mathcal{F}_e -tangle. So instead, there exists an S_k -tree over \mathcal{F}_e .

Among all S_k -trees over \mathcal{F}_e pick an irredundant one, (T, α) say, whose associated separations efficiently distinguishes as many pairs of profiles as possible. Let us suppose that some pair of profiles P_1, P_2 is not distinguished efficiently by this tree.

Consider the nodes v_{P_1}, v_{P_2} of this tree corresponding to P_1 and P_2 . These nodes are distinct, since every star in \mathcal{F}_e is contained in at most one profile. Moreover, we can assume without loss of generality, that in no node on the path between v_{P_1} and v_{P_2} there lives a profile Q : In that case either the pair P_1, Q or the pair Q, P_2 would not be efficiently distinguished by (T, α) either, so we could consider them instead.

Let \vec{s}_{P_1} be the separation associated to the first edge on the path from v_{P_1} to v_{P_2} and let \vec{s}_{P_2} be the separation associated to the first edge on the path from v_{P_2} to v_{P_1} . There exists a separation t which efficiently distinguishes P_1 and P_2 and is nested with \vec{s}_{P_1} and \vec{s}_{P_2} : if $\vec{t} \in P_1$ is not nested with, say \vec{s}_{P_1} , we know by property $\text{Eff}(P)$ that $\vec{s}_{P_1} \vee \vec{t}$ needs to have order at least $|s_{P_1}|$, thus $\vec{s}_{P_1} \wedge \vec{t}$ has order at most $|t|$, so it efficiently distinguishes P_1 and P_2 and is nested with \vec{s}_P . Thus by the fish Lemma 2.1, there indeed needs to exist such a t which efficiently distinguishes P_1 and P_2 and is nested with \vec{s}_{P_1} and \vec{s}_{P_2} . Moreover, t has an orientation such that $\vec{s}_{P_1} \leq \vec{t} \leq \vec{s}_{P_2}$, otherwise the existence of t again contradicts either property $\text{Eff}(P)$ or $\text{Eff}(Q)$. Note that \vec{t} thus is a splice between \vec{s}_{P_1} and \vec{s}_{P_2} and therefore \vec{t} emulates \vec{s}_{P_1} for \mathcal{F}_e and \vec{t} emulates \vec{s}_{P_2} for \mathcal{F}_e .

Let T_{P_1} be the subtree of T consisting of the component of $T - v_{P_1}$ which contains v_{P_2} together with v_{P_1} and similarly let T_{P_2} be the subtree consisting of the component of $T - v_{P_2}$ containing v_{P_1} together with v_{P_2} .

We consider the trees (T_{P_1}, α_{P_1}) and (T_{P_2}, α_{P_2}) obtained from $(T_{P_1}, \alpha \upharpoonright T_{P_1})$ and $(T_{P_2}, \alpha \upharpoonright T_{P_2})$ by applying the shifts $f \downarrow_{\vec{t}}^{\vec{s}_{P_1}}$ and $f \downarrow_{\vec{t}}^{\vec{s}_{P_2}}$, respectively. Consider now the tree (T', α') obtained from these two trees by identifying the respective edges associated with \vec{t} . By applying Lemma 2.8 with the two shifted trees the combined tree is again over \mathcal{F}_e . We may again assume it to be irredundant. We are going to show that it efficiently distinguishes more pairs of profiles than (T, α) .

Let Q_1, Q_2 be a pair profiles which were efficiently distinguished by a separation \vec{r} associated to an edge of (T, α) . If \vec{r} is not associated to any

edge of (T', α') , then, without loss of generality, either $\vec{s}_{P_1} \leq \vec{r} \leq \vec{s}_{P_2}$ or both $\vec{s}_{P_1} \leq \vec{r}$ and $\vec{s}_{P_2} \leq \vec{r}$.

In the first case \vec{r} distinguishes P_1 and P_2 and therefore $|\vec{r}| > |\vec{t}|$. By the definition of the shift, our tree (T', α') contains both, $\vec{r} \vee \vec{t}$ and $\vec{r} \wedge \vec{t}$, and both of them have order at most the order of \vec{r} , by Lemma 2.9. However, one of $\vec{r} \vee \vec{t}, \vec{r} \wedge \vec{t}$ and \vec{t} distinguishes Q_1 and Q_2 and does so efficiently.

In the second case, by the definition of the shift, our tree (T', α') contains both, $\vec{r} \vee \vec{t}$ and $\vec{r} \vee \vec{t}$, and both of them have order at most the order of \vec{r} , again by Lemma 2.9. Again, one of $\vec{r} \vee \vec{t}$ and $\vec{r} \vee \vec{t}$ distinguishes Q_1 and Q_2 and does so efficiently.

Thus, since (T', α') additionally efficiently distinguishes P_1 and P_2 with t , this contradicts the choice of (T, α) . \square

6. Degrees in efficient trees of tangles

In this section we apply our method from Section 4 to Theorem 5.1 to obtain a tree of tangles of low degree, but this time one which efficiently distinguishes the profiles. That is, we are interested in the minimal degrees of a tree of tangles whose associated separations efficiently distinguish all regular profiles of S_k .

Extending the definitions of Section 4, let us say that a tree of tangles (T, α) for S_k is *efficient*, if the set of edge labels not only distinguishes all regular profiles of S_k , but does so efficiently.

Given a k -profile P , we denote by $\delta_e(P)$ the minimal size of a star $\sigma \subseteq P$ with property $\text{Eff}(P)$ which distinguishes P from all other regular profiles of S_k , i.e., every other regular profile orients some $\vec{s} \in \sigma$ as \vec{s} . Note that, by Lemma 5.4, there exists such a star for every regular profile P , thus $\delta_e(P)$ is a well defined natural number.

We denote by $\delta_{e,\max}$ the maximum of $\delta_e(P)$ over all regular profiles P . We can give a bound on $\delta_e(P)$ which is not in terms of stars or nested sets:

Lemma 6.1. *Let P be a regular k -profile in U and let $D_P \subseteq P$ be a subset of P which contains, for every regular k -profile $P' \neq P$ in U , a separation which efficiently distinguishes P from P' . Let us denote as m the number of maximal elements of D_P . Then $\delta_e(P) \leq m$.*

Proof. It is enough to consider a set $D_P \subseteq P$ such that $m = |\max D_P|$ is as small as possible. Moreover, we may assume without loss of generality that every element of D_P distinguishes P efficiently from some other profile in \mathcal{P} , since we could otherwise remove it from D_P . We may furthermore assume that, subject to all this, D_P is chosen so that $\max D_P$ is \leq -minimal.

Furthermore we may suppose that, for separations $\vec{r} \leq \vec{s}$ in D_P , the order of \vec{r} is lower than the order of \vec{s} , since otherwise we could just remove \vec{r} from D_P .

If the maximal separations in D_P are pairwise nested, they satisfy property [Eff\(\$P\$ \)](#) by the fact that they distinguish P efficiently from some other profile P' . Further, every profile P' is distinguished from P by some maximal separation in D_P : there is an efficient P - P' distinguisher $\vec{s} \in D_P$ and thus a maximal separation $\vec{t} \geq \vec{s}$ in D_P also distinguishes P from P' . Hence, if the maximal elements of D_P are pairwise nested, they are a candidate for $\delta_e(P)$ and therefore witness that $\delta_e(P) \leq m$.

So suppose that this is not the case, so two maximal separations $\vec{s}, \vec{t} \in D_P$ cross and, without loss of generality, $|\vec{s}| \leq |\vec{t}|$. By the definition of D_P , there is a profile P_s which is efficiently distinguished from P by $\vec{s} \in D_P$. Similarly, there is such a profile P_t for \vec{t} .

Since D_P was chosen to have as few maximal elements as possible, the separation $\vec{s} \vee \vec{t}$ has greater order than t : otherwise we could, by consistency and the profile property, replace \vec{t} in D_P by $\vec{s} \vee \vec{t}$. Thus, by submodularity, the order of $\vec{s} \wedge \vec{t}$ is less than the order of \vec{s} . In particular, by efficiency of s and t , neither P_s nor P_t contains $(\vec{s} \wedge \vec{t})^* = \vec{s} \vee \vec{t}$.

Thus $\vec{s} \wedge \vec{t}$ and $\vec{s} \vee \vec{t}$ have order precisely $|\vec{s}|$ and $|\vec{t}|$, respectively: if one of them had lower order this would, by the profile property, contradict the fact that s or t , respectively, efficiently distinguishes P from P_s or P_t , respectively. This means that, in particular, $\vec{s} \wedge \vec{t}$ efficiently distinguishes P from P_s .

For every $\vec{r} \leq \vec{s}$ in D_P we have assumed $|\vec{r}| < |\vec{s}|$. Both $\vec{r} \wedge \vec{t}$ and $\vec{r} \vee \vec{t}$ have at most the order of \vec{r} due to submodularity, the efficiency of \vec{t} , the profile property and consistency, analogue to the above.

Let us consider the set D'_P obtained from D_P by removing all $\vec{r} \leq \vec{s}$, and adding $\vec{s} \wedge \vec{t}$ as well as, for every $\vec{r} \leq \vec{s}$, any $\vec{r} \wedge \vec{t}$ and $\vec{r} \vee \vec{t}$ which efficiently distinguishes P from some other profile. By the above, this set D'_P distinguishes P from every other regular profile, and is a candidate for D_P . The maximal separations of D'_P and of D_P are the same except that \vec{s} in D_P is replaced by $\vec{s} \wedge \vec{t}$ in D'_P . This contradicts the choice of D_P with \leq -minimal maximal elements. \square

To limit the degree of the node of P in our tree of tangles we want to remove from \mathcal{F}_e all the stars which are contained in P but are larger than $\delta_e(P)$. In order to achieve a maximum degree of $\delta_{e,\max}$ we also need to limit the size of the stars in \mathcal{F}_e which are contained in no profile to $\delta_{e,\max}$. As in Section 4, we cannot limit the maximum degrees below 3. Along the lines

of the proof of Lemma 3.2, the next lemma shows that we can find, in every consistent orientation O of \vec{S}_k which is not a profile, a star of size 3 contained in O and in no profile.

Lemma 6.2. *Every consistent orientation O of \vec{S}_k which is not a profile contains a star σ of size 3 which is not contained in any profile.*

Proof. As O is not a profile, there are $\vec{s}, \vec{t} \in O$ such that $\vec{s} \wedge \vec{t} \in O$. By submodularity, either $\vec{s} \wedge \vec{t}$ or $\vec{s} \wedge \vec{t} \in \vec{S}$, let us suppose the former one. Then $\sigma = \{\vec{s} \wedge \vec{t}, \vec{t}, \vec{s} \wedge \vec{t}\}$ is a star in O and σ cannot be contained in any profile: any profile P needs to contain either \vec{s} or \vec{s} , and the profile property implies that P then cannot contain both, $\vec{s} \wedge \vec{t}$ and $\vec{s} \wedge \vec{t}$. \square

We can now show the following variant of Theorem 5.1, which shows that we can find a tree of tangles of bounded degree:

Theorem 6.3. *Let \vec{U} be a submodular universe and let \mathcal{P} be the set of regular profiles of \vec{S}_k . Then there exists tree of tangles (T, α) such that, for every profile $P \in \mathcal{P}$, the degree of P in (T, α) is $\delta_e(P)$ and the maximal degree of T is at most $\max\{\delta_e(\mathcal{P}), 3\}$.*

Proof. Let \mathcal{F}_e^s be the subset of \mathcal{F}_e consisting of, for every profile P , all stars from \mathcal{F}_e of size $\delta_e(P)$ contained in P , together with all stars of size at most $\max\{\delta_e(\mathcal{P}), 3\}$ from \mathcal{F}_e not contained in any profile. For any star σ and any shift $\sigma_{\vec{s}}$ of σ we have $|\sigma| \geq |\sigma_{\vec{s}}|$. Further, S_k is \mathcal{F}_e -separable by Corollary 5.3. Moreover, the shift of a star cannot contain any profile which does not contain the original star by Lemma 3.3, thus S_k is also \mathcal{F}_e^s -separable.

Thus, all we need to show is that applying Theorem 1.2 cannot result in an \mathcal{F}_e^s -tangle, the rest of the proof can then be carried out as the proof of Theorem 5.1: Instead of S -trees over \mathcal{F}_e we now consider S -trees over \mathcal{F}_e^s , and observe that the shifting argument in the proof of Theorem 5.1 again shifts stars in \mathcal{F}_e^s to stars in \mathcal{F}_e^s .

However, applying Theorem 1.2 indeed cannot result in an \mathcal{F}_e^s -tangle: Such a tangle cannot be a regular profile, since by our definition of $\delta_e(P)$, there is a star in \mathcal{F}_e^s contained in P . But every consistent orientation which is not a regular profile either contains a star $\{\vec{s}\}$ for a cosmall separation \vec{s} – each such star is also contained in \mathcal{F}_e – or contains, by Lemma 6.2 a star of size 3 not contained in any profile. Either such star is also contained in \mathcal{F}_e^s by definition. \square

7. Tangles of mixed orders

In this section we would like to use the ideas from Section 5 to obtain a proof of Theorem 2.2 using tangle-tree duality. The challenge of Theorem 2.2 compared to Theorem 5.1 is that the set of profiles \mathcal{P} considered in Theorem 2.2 consists of profiles of different orders. In particular, there might be profiles P_1 and P_2 in \mathcal{P} which are efficiently distinguished by separations of order k , say, and there might be another profile $Q \in \mathcal{P}$ which has only order $l < k$ and thus does not orient the separations which efficiently distinguish P_1 and P_2 . Thus, we cannot simply require the stars in our set \mathcal{F} to be contained in at most one profile: the resulting S -tree over \mathcal{F} would not necessarily distinguish all profiles in \mathcal{P} , for example it might not distinguish the profiles P_1 and Q from above. Our solution to this problem will be to restrict the set of stars further by additionally requiring that all the separations in a star in \mathcal{F} ‘could be oriented’ by every profile in \mathcal{P} , even if that profile has lower order than the separation considered.

With this further restricted set of stars however S will no longer be \mathcal{F} -separable, but it will only fail to do so under rather specific circumstances. Thus in order to obtain a result in the fashion of Theorem 2.2, we shall first prove a slightly stronger version of Theorem 1.2, which allows us to exclude this specific situation in the requirement of \mathcal{F} -separability. The proof of this stronger version of Theorem 1.2 is different from the original one of Theorem 1.2 in [7]: our proof of Theorem 7.1 also serves as an alternative proof of Theorem 1.2 which is slightly shorter than the original one and perhaps neater and less technical.

7.1. A short adventure into tangle-tree duality

As mentioned above, we will prove the following slight strengthening of Theorem 1.2:

Theorem 7.1. *Let U be a finite universe, $S \subseteq U$ a separation system, and $\mathcal{F} \subseteq 2^{\bar{S}}$ a set of stars such that \mathcal{F} is standard for S and S is critically \mathcal{F} -separable. Then precisely one of the following holds:*

- *there is an S -tree over \mathcal{F} ;*
- *there is an \mathcal{F} -tangle of S .*

This theorem is a strengthening in the sense that, we weaken the technical assumption that S be \mathcal{F} -separable to only require \mathcal{F} -separability for those separations whose inverse lies in no star of \mathcal{F} , rather than for all separations in \bar{S} . Formally:

A separation \vec{r} in \vec{S} is \mathcal{F} -critical if $\vec{r} \in \sigma$ for some $\sigma \in \mathcal{F}$, but there is no $\sigma' \in \mathcal{F}$ with $\sigma' \cap r = \{\vec{r}\}$. Observe that if $\vec{r} \in \vec{S}$ is \mathcal{F} -critical, then \vec{r} is non-degenerate and not forced by \mathcal{F} , and in particular \vec{r} is nontrivial in S since \mathcal{F} is standard for S . We say that S is *critically \mathcal{F} -separable* if for all \mathcal{F} -critical $\vec{r}, \vec{r}' \in \vec{S}$ with $\vec{r} \leq \vec{r}'$ there exists an $s_0 \in S$ with an orientation \vec{s}_0 that emulates \vec{r} in \vec{S} for \mathcal{F} and such that \vec{s}_0 emulates \vec{r}' in \vec{S} for \mathcal{F} . Clearly, if S is \mathcal{F} -separable, then S is critically \mathcal{F} -separable.

The core argument of our proof of Theorem 7.1 is the following: if there is no S -tree over \mathcal{F} , then every ‘attempt’ at such an S -tree must fail. Thus, if one starts with some star σ in \mathcal{F} as the basis for such an S -tree, and then ‘glues’ for each $\vec{s} \in \sigma$ with $\{\vec{s}\} \notin \mathcal{F}$ some star from \mathcal{F} onto \vec{s} that contains \vec{s} , then one must at some point be unable to find such a star. The resulting attempt is then an S -tree that is ‘over \mathcal{F} ’ only for all internal vertices, but not necessarily at the leaves. The fundamental strategy of the proof presented here is to collect the set of all these leaf-separations at which the S -tree attempts get stuck, and then turn this set into the basis of an \mathcal{F} -tangle.

The strategy described here is already present in Mazoit’s proof ([12]) of the classical duality theorem for brambles and tree-width in graphs. In [1] Diestel gives a proof of this graph-theoretic duality theorem that is derived from his and Oum’s original proof of Theorem 1.2, applied to the specific \mathcal{F} corresponding to tree decompositions of a certain width. Curiously Mazoit’s and Diestel’s graph-theoretic proofs are quite similar. One could therefore argue that the strategy for the proofs here comes from re-translating Diestel’s translation of Theorem 1.2 to this specific graph application back into abstract separation systems.

Let us now make this sketch of a proof formal:

Proof of Theorem 7.1. Let U be a finite universe, $S \subseteq U$ a separation system, and $\mathcal{F} \subseteq 2^{\vec{S}}$ a set of stars such that \mathcal{F} is standard for S and that S is critically \mathcal{F} -separable.. We may assume that $\emptyset \notin \mathcal{F}$.

We shall need the following definitions. An S -tree attempt is an S -tree (T, α) with at least one edge and $\alpha(t) \in \mathcal{F}$ for every internal node t of T . For a leaf t of an S -tree attempt (T, α) the *incoming label* of t is the separation \vec{r} for which $\{\vec{r}\}$ is associated with t in (T, α) ; the *outgoing label* of t is then \vec{r} . We call such an \vec{r} a *petal* of (T, α) if $\{\vec{r}\} \notin \mathcal{F}$.

We will prove the following assertion which is equivalent to the tangle-tree duality theorem:

(*) S has an \mathcal{F} -tangle if and only if every S -tree attempt has a petal.

Since the S -tree attempts without petals are exactly the S -trees that are over \mathcal{F} , $(*)$ immediately implies Theorem 1.2.

For our proof of $(*)$ we will use Lemma 2.3 to 2.6. These lemmas, roughly speaking, say that an S -tree over a set of stars may be assumed to be ‘cleaned up’, i.e. tight and irredundant; that such a cleaned up S -tree (T, α) is order-respecting; and that then each nontrivial separation can only appear once as a label. In short, cleaning up a given S -tree over stars enables us to apply Lemma 2.8 to it.

The ‘only if’ direction of $(*)$ is clear, so let us show the backward direction.

Let $\vec{L} \subseteq \vec{S}$ be the set of all petals of S -tree attempts. Suppose first that we can find some $P \subseteq \vec{L}$ which is a consistent and antisymmetric set that contains at least one petal of every S -tree attempt. Then P is an \mathcal{F} -tangle of S : Indeed, P orients each separation s in S since it contains a petal of the S -tree attempt that is just a single edge labeled with s ; and P avoids each star σ in \mathcal{F} since it is antisymmetric and contains a petal of the S -tree attempt consisting of one internal node for σ and one leaf for every element of σ .

So let us show that we can find such a set $P \subseteq \vec{L}$. For this pick a $P \subseteq \vec{L}$ that is minimal with respect to inclusion subject to the conditions that P contains at least one petal of every S -tree attempt and is down-closed in \vec{L} , that is, so that $\vec{p} \in P$ for all $\vec{p} \in \vec{L}$ with $\vec{p} \leq \vec{q}$ for some $\vec{q} \in P$. Such a P exists since \vec{L} itself is a candidate. We claim that this P is antisymmetric and consistent.

So suppose that P is not antisymmetric, or not consistent. Then there are $\vec{r} \neq \vec{s}$ in P with $\vec{r} \leq \vec{s}$. In particular we can take \vec{r} and \vec{s} to be maximal elements of P . Neither \vec{r} nor \vec{s} can be co-trivial in S since \mathcal{F} is standard for S and \vec{r} and \vec{s} are petals. Therefore neither of the two can be trivial or degenerate either, since this would imply that the other one is co-trivial.

By picking \vec{r} and \vec{s} among the maximal elements of P we ensure that both $P \setminus \{\vec{r}\}$ and $P \setminus \{\vec{s}\}$ are still down-closed in \vec{L} . Thus, by the minimality of P , there are S -tree attempts (T_r, α_r) and (T_s, α_s) whose only petals that lie in P are \vec{r} and \vec{s} , respectively. We may assume (T_r, α_r) and (T_s, α_s) to be tight and irredundant by Lemma 2.5 and 2.6, which implies that \vec{r} and \vec{s} are the incoming label of exactly one leaf of T_r and T_s , respectively, by Lemma 2.4.

We claim that \vec{r} is \mathcal{F} -critical. To see this, let v_r be the leaf of T_r whose incoming edge is labeled with \vec{r} , and let w_r be the neighbour of v_r in T_r . Then w_r has an incoming edge labelled with \vec{r} , and we must have $\alpha_r(w_r) \in \mathcal{F}$, witnessing that \vec{r} lies in some star in \mathcal{F} : for if $\alpha_r(w_r) \notin \mathcal{F}$, then w_r would

be a leaf of T_r and \vec{r} a petal of (T_r, α_r) . By $\vec{r} \leq \vec{s}$ and P being down-closed in \vec{L} we would then have $\vec{r} \in P$, contrary to the assumption that \vec{r} is the only petal of (T_r, α_r) which lies in P . Suppose now that $\sigma \cap \{\vec{r}, \vec{r}'\} = \{\vec{r}'\}$ for some $\sigma \in \mathcal{F}$. Then we can extend (T_r, α_r) by σ at the leaf at which \vec{r}' appears. This extension of (T_r, α_r) is then an S -tree attempt, of which P must contain a petal. Since $\vec{r}' \notin \sigma$ this petal would be strictly larger than \vec{r} , contradicting the maximality of \vec{r} in P .

A similar argument shows that \vec{s} is \mathcal{F} -critical. By the assumption that S is critically \mathcal{F} -separable we thus find an $s_0 \in S$ with an orientation \vec{s}_0 that emulates \vec{r} in \vec{S} for \mathcal{F} and such that \vec{s}_0 emulates \vec{s} in \vec{S} for \mathcal{F} . Let v_r and v_s be the leaves of T_r and T_s with incoming labels \vec{r} and \vec{s} , respectively. Set $\alpha'_r := (\alpha_r)_{v_r, \vec{s}_0}$ and $\alpha'_s := (\alpha_s)_{v_s, \vec{s}_0}$. Then Lemma 2.8 says that (T_r, α'_r) and (T_s, α'_s) are S -tree attempts in which v_r is the unique leaf of T_r with incoming label \vec{s}_0 , and that v_s is the unique leaf of T_s with incoming label \vec{s}_0 . Let (T, α) be the S -tree obtained from (T_r, α'_r) and (T_s, α'_s) by identifying $v_r \in T_r$ with the neighbour of v_s in T_s and vice-versa, and extending the maps α'_r and α'_s accordingly.

This (T, α) , too, is an S -tree attempt since every internal node of (T, α) corresponds to an internal node of (T_r, α'_r) or (T_s, α'_s) .

We claim that P , contrary to its definition, contains no petal of (T, α) . To see this we consider the leaves of (T, α) and note that by definition every leaf of T corresponds to either a leaf of T_r other than v_r , or to a leaf of T_s other than v_s . Thus, let us first consider some leaf t of T_r other than v_r , and let \vec{p} be the incoming label of t in (T_r, α_r) . Then $\vec{r} \leq \vec{p}$ by Lemma 2.3, and thus the incoming label of t in (T, α) is $\vec{s}_0 \vee \vec{p}$. If $\alpha_r(t) = \{\vec{p}\} \in \mathcal{F}$, then $\alpha(t) = \{\vec{s}_0 \vee \vec{p}\} \in \mathcal{F}$ since \vec{s}_0 emulates \vec{r} in \vec{S} for \mathcal{F} . Consequently, if $\vec{s}_0 \vee \vec{p}$ is a petal of (T, α) , then \vec{p} is a petal of (T_r, α_r) . In particular, since $\vec{p} \neq \vec{r}$ and \vec{r} is the only petal of P from (T_r, α_r) , and P is down-closed in \vec{L} , we know that $(\vec{s}_0 \vee \vec{p}) \geq \vec{p}$ cannot be a petal of (T, α) that lies in P .

By the same argument those leaves of T which are leaves of T_s other than v_s cannot give rise to a petal of (T, α) in P , either. Hence P contains no petal from (T, α) , causing a contradiction. This finishes the proof that P is consistent and antisymmetric and hence an \mathcal{F} -tangle of S . \square

7.2. Obtaining a tree-of-tangles theorem for different order tangles from tangle-tree duality

Theorem 7.1 now allows us to use our methods from Theorem 5.1 to prove a tree-of-tangles theorem for different order tangles. More specifically we will obtain a result similar to Theorem 2.2, however our construction only works

in distributive universes – that is, $\vec{r} \vee (\vec{s} \wedge \vec{t}) = (\vec{r} \vee \vec{s}) \wedge (\vec{r} \vee \vec{t})$, always – since we need the following result from [8], which can be also found in [5]:

Lemma 7.2 ([5, Theorem 3.11], [8, Theorem 1], strong profile property). *Let \vec{U} be a distributive universe and $\vec{S} \subseteq \vec{U}$ structurally submodular, then for any profile P of S and any \vec{r} and $\vec{s} \in P$ there does not exist any $\vec{t} \in P$ such that $\vec{r} \vee \vec{s} \leq \vec{t}$.*

Moreover, our method will not allow us to distinguish all robust profiles, instead we need a slight strengthening of robustness: We say that a k -profile P is *strongly robust*, if for any $\vec{s} \in P$ and $\vec{r} \in \vec{U}$ where $\vec{s} \vee \vec{r}$ and $\vec{s} \vee \vec{r}$ both have at most the order of \vec{s} one of $\vec{s} \vee \vec{r}$ and $\vec{s} \vee \vec{r}$ is in P . Note that most instances of tangles, for example tangles in graphs, are strongly robust profiles.

For this section let \vec{U} be a distributive submodular universe and let \mathcal{P} be some set of pairwise distinguishable strongly robust profiles in \vec{U} (possibly of different order).

To handle the issue, that not all separations in a tree-of-tangles for profiles of different orders are oriented by all the considered profiles, we introduce the following additional definition: A consistent orientation O of S_k *weakly orients a separation s as \vec{s}* if O contains a separation \vec{r} such that $\vec{s} \leq \vec{r}$. If we want to omit s we just say O *weakly contains \vec{s}* .

We will now only consider stars of separations where every separation is at least weakly oriented by all the profiles in \mathcal{P} . Specifically, we work with the set \mathcal{F}_d consisting of all stars σ with the following properties:

1. There exists at most one profile $P \in \mathcal{P}$ such that $\sigma \subseteq P$.
2. For every profile $P \in \mathcal{P}$ such that $\sigma \not\subseteq P$ there exists $\vec{s} \in \sigma$ such that P weakly orients s as \vec{s} .
3. If there exists a $P \in \mathcal{P}$ such that $\sigma \subseteq P$, then σ satisfies property [Eff\(\$P\$ \)](#).

We want to show that U is critically \mathcal{F}_d -separable, and our first step to do so is to show that splices – which we want to use in separability – are weakly oriented by every profile in \mathcal{P} .


Lemma 7.3. *Let \vec{U} be a distributive submodular universe and let \mathcal{P} be a set of strongly robust profiles in \vec{U} . Suppose that \vec{r} and \vec{s} are \mathcal{F}_d -critical separations in \vec{U} with $\vec{r} \leq \vec{s}$, then every splice between \vec{r} and \vec{s} is weakly oriented by every profile in \mathcal{P} .*

Proof. Since \vec{r} and \vec{s} are \mathcal{F}_d -critical, they are contained in some star in \mathcal{F}_d and hence weakly oriented by every profile in \mathcal{P} .

Let t be a splice between \vec{r} and \vec{s} . If t is not weakly oriented by every profile in \mathcal{P} , then \mathcal{P} contains a profile P of order at most $|t|$ which weakly orients r as \vec{r} and s as \vec{s} , since every witnessing separation that a profile weakly orients r as \vec{r} or s as \vec{s} also witnesses that it weakly orients t . Let M_r^P be the set of all separations \vec{w}_r in P satisfying $\vec{r} \leq \vec{w}_r$ and having minimal possible order with that property. Let $\vec{w}_r \in M_r^P$ be chosen \leq -maximally. Let \vec{w}_s be defined for s , accordingly.

Observe that if $\vec{w}_r \leq \vec{s}$, respectively, then, by the order-minimality of M_r^P , the order of \vec{w}_r is at least $|t|$ so P orients t , which contradicts the assumption that P does not weakly orient t . Similarly, $\vec{w}_s \leq \vec{r}$ results in a contradiction.

Suppose now that \vec{w}_r crosses \vec{s} .



We claim that every profile P' in \mathcal{P} which weakly orients s as \vec{s} also weakly contains either $\vec{s} \vee \vec{w}_r$ or $\vec{s} \vee \vec{w}_s$. This then implies that the set $\{\vec{s}, \vec{s} \wedge \vec{w}_r, \vec{s} \wedge \vec{w}_s\}$ is a star in \mathcal{F}_d , which will contradict the \mathcal{F}_d -criticality of \vec{s} .

So suppose that P' weakly orients s as \vec{s} , witnessed by some $\vec{w} \in P'$ with $\vec{w} \geq \vec{s}$.


If $\vec{w}_r \vee \vec{w}$ had order at most the order of \vec{w}_r , this would contradict the choice of \vec{w}_r : By Lemma 7.2 applied to the separations $\vec{w}_r, \vec{w}_s, \vec{w} \wedge \vec{w}_r$, the profile P would need to contain $\vec{w}_r \vee \vec{w}$ which contradicts the choice of \vec{w}_r being \leq -maximal in M_r^P .

Similarly, if $\vec{w} \wedge \vec{w}_r$ had order less than the order of \vec{w}_r , this would contradict the choice of \vec{w}_r : By consistency P would need to contain $\vec{w} \wedge \vec{w}_r$ which contradicts the definition of M_r^P , from which \vec{w}_r was chosen.

Thus, by submodularity, $\vec{w} \wedge \vec{w}_r$ has order less than the order of w , and $\vec{w} \wedge \vec{w}_s$ has order at most the order of w . Hence, as P' is strongly robust, P'

contains either $\vec{w} \vee \vec{w}_r$ or $\vec{w} \vee \vec{w}_r$ and therefore either weakly orients $\vec{s} \wedge \vec{w}_r$ as $\vec{s} \vee \vec{w}_r$ or $\vec{s} \wedge \vec{w}_r$ as $\vec{s} \vee \vec{w}_r$.

This proves the claim which results in a contradiction to the assumption that \vec{s} is \mathcal{F}_d critical. Thus we may suppose that \vec{w}_r does not cross \vec{s} and, by a symmetric argument, that \vec{w}_s does not cross \vec{r} . Hence $\vec{r} \leq \vec{w}_s$ and $\vec{s} \leq \vec{w}_r$. We may therefore assume without loss of generality that $\vec{w}_r = \vec{w}_s$.



If $\vec{w}_r = \vec{w}_s$ crosses t , then, by the choice of t , that neither $\vec{w}_r \wedge \vec{t}$ nor $\vec{w}_r \wedge \vec{t}$ has order less than $|t|$, thus $\vec{w}_r \vee \vec{t}$ and $\vec{w}_r \vee \vec{t}$ both have order at most the order of \vec{w}_r . By the strong robustness of P applied to \vec{w}_r , $\vec{w}_r \wedge \vec{t}$ and $\vec{w}_r \wedge \vec{t}$, we know that either $\vec{w}_r \vee \vec{t} \in P$ or $\vec{w}_r \vee \vec{t} \in P$. However, both contradict the \leq -maximal choice of \vec{w}_r . So, instead \vec{w}_r is nested with t , that is, t has an orientation \vec{t} such that $\vec{t} \leq \vec{w}_r$, so t is weakly oriented by P , as claimed. \square

Note that the assumption that our profiles are *strongly* robust is essential in this argument, for example for the case $\vec{w}_r = \vec{w}_s$: If we only assume robustness, we can not conclude that P contains either $\vec{w}_r \vee \vec{t}$ or $\vec{w}_r \vee \vec{t}$ and thus would not obtain a contradiction.

The next step is to verify that shifting with a splice as in Lemma 7.3 maps stars in \mathcal{F}_d to stars in \mathcal{F}_d , which will prove that U is critically \mathcal{F}_d -separable:

Lemma 7.4. *Let r and s_0 be separations which are weakly oriented by every profile in \mathcal{P} and suppose that \vec{s}_0 is a splice for \vec{r} . Let $\sigma \in \mathcal{F}_d$ be a star which contains a separation $\vec{x} \geq \vec{r}$. Then the shift $\sigma_{\vec{x}}^{\vec{s}_0}$ of σ from \vec{x} to \vec{s}_0 is again an element of \mathcal{F}_d .*

Proof. Since \vec{s}_0 is a splice for \vec{r} , by Lemma 2.9, $\vec{s} \vee \vec{s}_0$ has at most the order of \vec{s} for every $\vec{s} \geq \vec{r}$.

Let σ be any star in \mathcal{F}_d containing a separation $\vec{x} \geq \vec{r}$. By the above, if $\sigma \subseteq \vec{S}_k$ for some k , then also the shift $\sigma_{\vec{x}}^{\vec{s}_0}$ is a subset of \vec{S}_k . Hence by Lemma 3.3, every profile in U which contains $\sigma_{\vec{x}}^{\vec{s}_0}$ also contains σ . Now if some profile P contains σ , then P orients every separation in $\sigma_{\vec{x}}^{\vec{s}_0}$, and thus either P contains the inverse of some separation in $\sigma_{\vec{x}}^{\vec{s}_0}$ or $\sigma_{\vec{x}}^{\vec{s}_0} \subseteq P$.

Hence, by Lemma 5.2 it is enough to show that every profile from \mathcal{P} which, for some $\vec{y} \in \sigma$, weakly contains \vec{y} also weakly contains \vec{y}' for some separation $\vec{y}' \in \sigma_{\vec{x}}^{\vec{s}_0}$.

So suppose such a profile P , for some $\vec{y} \in \sigma$, weakly contains \vec{y} and suppose that this is witnessed by $\vec{w}_y \in P$. If $\vec{r} \leq \vec{y}$, then \vec{y} is shifted onto $\vec{y} \wedge \vec{s}_0$ and therefore \vec{w}_y also witnesses that P weakly contains $\vec{y} \wedge \vec{s}_0$ while $\vec{y} \vee \vec{s}_0 \in \sigma_{\vec{x}}^{\vec{s}_0}$. Thus we may suppose that $\vec{r} \leq \vec{y}$ and therefore that \vec{y} is shifted onto $\vec{y} \vee \vec{s}_0$.

If P weakly orients s_0 as \vec{s}_0 , then P also weakly contains $\vec{y} \wedge \vec{s}_0 \leq \vec{s}_0$ while $\vec{y} \vee \vec{s}_0 \in \sigma_{\vec{x}}^{\vec{s}_0}$.

Thus we may suppose that P weakly orients s_0 as \vec{s}_0 , witnessed by $\vec{w}_0 \in P$.

By our assumptions on s_0 we know that the order of $\vec{s}_0 \wedge \vec{w}_y$ is at least the order of s_0 and thus, by submodularity, $\vec{s}_0 \wedge \vec{w}_y$ has order at most the order of w_y , i.e., it is oriented by P . By Lemma 7.2 applied to $\vec{w}_0, \vec{w}_y \in P$ and $\vec{s}_0 \wedge \vec{w}_y$ we can therefore conclude that P contains $\vec{s}_0 \vee \vec{w}_y$, i.e., P weakly contains $\vec{y} \vee \vec{s}_0 \leq \vec{s}_0 \vee \vec{w}_y$. \square

In order to use our stronger tangle-tree duality theorem Theorem 7.1 with our set \mathcal{F}_d of stars to obtain a tree of tangles for strongly robust profiles it only remains for us to show that this application cannot result in an \mathcal{F}_d -tangle. We do this in the following two lemmas.

Lemma 7.5. *For every profile P in \vec{U} and every set \mathcal{P}' of strongly robust profiles in \vec{U} distinguishable from P , there exists a nested set N which distinguishes P efficiently from all the profiles in \mathcal{P}' .*

Proof. For every profile $Q \in \mathcal{P}'$ pick a \leq -minimal separation $\vec{s}_Q \in P$ which efficiently distinguishes Q from P . We claim that the set N consisting of all these separations s_Q is nested and therefore as claimed.

So suppose that this is not the case, so \vec{s}_Q and $\vec{s}_{Q'}$, say, cross. We may assume without loss of generality that $|\vec{s}_Q| \leq |\vec{s}_{Q'}|$. Now $\vec{s}_Q \vee \vec{s}_{Q'}$ has order at least the order of $\vec{s}_{Q'}$ since otherwise, by the profile property, $\vec{s}_Q \vee \vec{s}_{Q'}$ would also distinguish P and Q' and would thus contradict the fact that $\vec{s}_{Q'}$ did so efficiently. Thus $|\vec{s}_Q \wedge \vec{s}_{Q'}| \leq |\vec{s}_Q|$.

Now Q' orients \vec{s}_Q and it cannot contain \vec{s}_Q since then, by the profile property, $\vec{s}_Q \wedge \vec{s}_{Q'}$ would also distinguish P and Q' efficiently and would therefore contradict the \leq -minimal choice of $\vec{s}_{Q'}$.

Thus $\vec{s}_Q \in Q'$. Now $|\vec{s}_Q \wedge \vec{s}_{Q'}| > |\vec{s}_{Q'}|$ since otherwise, again by the profile property, $\vec{s}_Q \wedge \vec{s}_{Q'}$ contradicts the \leq -minimal choice of $\vec{s}_{Q'}$.

Thus, by submodularity, $|\vec{s}_Q \wedge \vec{s}_{Q'}| < |\vec{s}_Q|$ and $|\vec{s}_Q \wedge \vec{s}_{Q'}| \leq |\vec{s}_Q|$. But, by strong robustness, either $\vec{s}_Q \vee \vec{s}_{Q'}$ or $\vec{s}_Q \wedge \vec{s}_{Q'}$ is in Q . In particular, $\vec{s}_Q \wedge \vec{s}_{Q'}$ or $\vec{s}_Q \vee \vec{s}_{Q'}$ efficiently distinguishes P and Q and therefore contradicts the \leq -minimal choice of \vec{s}_Q . \square

Unlike for structurally submodular separation systems in Lemma 3.5 or efficient distinguishers in Lemma 6.2, in this setup we can not necessarily find a star in \mathcal{F}_d which is contained in O but in no profile in \mathcal{P} for every orientation O of U which not include any profile in our set \mathcal{P} of strongly robust profiles. This is because we require that every profile in \mathcal{P} weakly orients a separation in our star outwards, but the stars constructed in Lemma 6.2, for example, do not necessarily have this property. Thus we are going to, instead, find a star σ contained in both O and exactly one profile from \mathcal{P} . Since each such star also lies in \mathcal{F}_d , this will be enough to ensure that our application of Theorem 7.1 does not result in an \mathcal{F}_d -tangle.

Lemma 7.6. *For every consistent orientation O of \vec{U} and every set $\mathcal{P} \neq \emptyset$ of distinguishable strongly robust profiles in \vec{U} there exists a star σ in \mathcal{F}_d contained in O .*

Proof. Pick a star σ (not necessarily from \mathcal{F}_d) with the following properties:

- (i) $\sigma \subseteq O$.
- (ii) σ is contained in at least one profile in \mathcal{P} .
- (iii) Property **Eff**(P) is satisfied for every profile $P \in \mathcal{P}$ such that $\sigma \subseteq P$.
- (iv) Every $P \in \mathcal{P}$ either contains σ or weakly contains \vec{s} for some separation $\vec{s} \in \sigma$.
- (v) For every separation $\vec{s} \in \sigma$ and any profile $\sigma \subseteq P$ there exists a profile $Q \in \mathcal{P}$ such that s is a efficient P - Q distinguisher.

Note that the empty set is such a star. Let us further assume that we choose our star σ fulfilling (i)–(v) so that as few profiles in \mathcal{P} as possible contain σ .

If only one profile contains σ , then $\sigma \in \mathcal{F}_d$ is as desired, so let us suppose for a contradiction that there are at least two such profiles.

Pick two such profiles $P_1, P_2 \supseteq \sigma$ such that the order of an efficient P_1 - P_2 -distinguisher is as small as possible. Pick an efficient P_1 - P_2 -distinguisher s which crosses as few elements of σ as possible. O orients s , say $\vec{s} \in O$. If s

is nested with σ , the maximal elements of $\sigma \cup \{\vec{s}\}$ form a star violating the definition of σ : Every profile containing this new star also contains σ . To see that (iii) is fulfilled, note that there is no profile $P \supseteq \sigma$ in \mathcal{P} such that $\vec{s} \in P$ for which there is a \vec{s}' of lower order than \vec{s} such that $\vec{s} \leq \vec{s}' \in P$, since such an \vec{s}' would be a distinguisher of lower order than \vec{s} for some pair of profiles containing σ , contrary to the choice of s .

Thus we may assume that s is not nested with σ , say s crosses $\vec{t} \in \sigma$. Since, by (v), there is some profile $Q \ni \vec{t}$ for which t is an efficient P_1 - Q -distinguisher, we know that at least one of $\vec{s} \wedge \vec{t}$ and $\vec{s} \vee \vec{t}$ has order at least the order of t : Otherwise this would contradict the fact that t is an efficient P_1 - Q -distinguisher by robustness (if $|t| < |s|$) or the profile property (if $|s| \leq |t|$) of Q .

Thus by submodularity the order of at least one of $\vec{s} \vee \vec{t}$ and $\vec{s} \wedge \vec{t}$ is at most the order of s and that separation is therefore also an efficient P_1 - P_2 -distinguisher (by the profile property and consistency), which would make it a better choice for s , a contradiction.

Thus σ contains precisely one profile and therefore, by construction, $\sigma \in \mathcal{F}_d$. □

Together with Theorem 7.1, these lemmas give a proof of a tree-of-tangles theorem for strongly robust profiles of different orders in a submodular universe. This theorem does not give efficient distinguishers; we will deal with efficiency in a later step.

Theorem 7.7 (Tree-of-tangles theorem for different orders). *Let $U = (\vec{U}, \leq, *, \vee, \wedge, | \cdot |)$ be a submodular distributive universe of separations. Then for every distinguishable set \mathcal{P} of strongly robust profiles in \vec{U} there is a nested set $T = T(\mathcal{P}) \subseteq U$ of separations such that:*

- (i) every two profiles in \mathcal{P} are distinguished by some separation in T ;
- (ii) for any profile $P \in \mathcal{P}$, any maximal $\vec{s} \in P \cap \vec{T}$ and any $\vec{s}' \in P$ such that $\vec{s} \leq \vec{s}'$ we have $|\vec{s}| \leq |\vec{s}'|$.

Proof. By Lemma 7.3 and Lemma 7.4 the set U is critically \mathcal{F}_d -separable for the set \mathcal{F}_d defined above. Thus we can apply Theorem 7.1. This can, by Lemma 7.6, not result in an \mathcal{F}_d -tangle, thus there is an U -tree over \mathcal{F}_d . By Lemma 2.5 we may assume this U -tree to be irredundant. The set of separations associated to edges of this tree is then a nested set T .

Every profile in \mathcal{P} induces a consistent orientation of T , since all the separations in T are weakly oriented by every profile in \mathcal{P} . The maximal elements of this orientation form a star σ_P in \mathcal{F}_d , and this star is a subset of P by the definition of \mathcal{F}_d .

To see that T distinguishes every pair of profiles in \mathcal{P} , consider two profiles P and Q in \mathcal{P} . These two profiles cannot induce the same orientation of T , since then $\sigma_P = \sigma_Q$ would be a subset of both P and Q , contradicting the definition of \mathcal{F}_d . Thus some $\vec{s} \in \sigma_P$ witnesses that P weakly orients some $\vec{t} \in \sigma_Q$ as \vec{t} and, vice versa, \vec{t} witnesses that Q weakly contains \vec{s} . Of these two separations s and t , the one of lower order is thus a P - Q -distinguisher in T .

Assertion (ii) is then immediate from the definition of \mathcal{F}_d . □

Note that the nested set constructed in Theorem 7.7 does not yet necessarily distinguish any two profiles *efficiently*. However, we can use Theorem 5.1 in combination with Theorem 7.7 to obtain such a set:

Theorem 7.8 (Efficient tree-of-tangles theorem for different order profiles). *Let $U = (\vec{U}, \leq, *, \vee, \wedge, | |)$ be a submodular distributive universe of separations. Then for every distinguishable set \mathcal{P} of strongly robust profiles in \vec{U} there is a nested set $T = T(\mathcal{P}) \subseteq U$ of separations such that every two profiles in \mathcal{P} are efficiently distinguished by some separation in T .*

Proof. Let k be the maximal order of a profile in U . Let T be the U -tree over \mathcal{F}_d from the proof of Theorem 7.7. We consider the \subseteq -maximal subtrees T_i of T with the property that no internal node of T_i corresponds to a profile in \mathcal{P} . Clearly $T = \bigcup_{i=1}^m T_i$ and no two T_i share an edge.

We are going to simultaneously replace each of the nested sets of separations corresponding to the T_i s with other separations in such a way that the resulting set of separations is still nested and we ensured that every pair of profiles contained in some T_i is efficiently distinguished by this new set of separations.

So, given some T_i , let \mathcal{P}_i be the set of profiles in \mathcal{P} living, in T , in one of the leaves of T_i . Let \vec{L}_i be the set of all separations associated to one of the directed edges adjacent and pointing away from such a leaf. Note that \vec{L}_i is a star. For every $\vec{s} \in \vec{L}_i$ let $P_s \in \mathcal{P}_i$ be the unique profile corresponding to a leaf of T_i and containing \vec{s} .

It is easy to check that for any two profiles P and Q in \mathcal{P}_i there is a efficient P - Q -distinguisher t which is nested with all of \vec{L}_i : Pick one t which is nested with as many separations from \vec{L}_i as possible. Now t cannot cross an $\vec{s} \in \vec{L}_i$ such that $P_s = P$ or $P_s = Q$, as in that case, for $\vec{t} \in P_s$, either $\vec{t} \vee \vec{s}$ or $\vec{t} \wedge \vec{s}$ would, by submodularity, consistency and the profile property, be an efficient P - Q -distinguisher and as such contradict the choice of t by Lemma 2.1. If on the other hand t crosses some $\vec{s} \in \vec{L}_i$, such that $P_s \notin \{P, Q\}$, then not both of $\vec{s} \vee \vec{t}$ and $\vec{s} \wedge \vec{t}$ can have order less than the

order of s by the profile property since, by (ii), there is no $\overleftarrow{s'} \in P_s$ such that $\overleftarrow{s} \leq \overleftarrow{s'}$ and $|\overleftarrow{s}| > |\overleftarrow{s'}|$. Thus the order of either $\overleftarrow{s} \vee \overleftarrow{t}$ or $\overleftarrow{s} \vee \overleftarrow{t}$ is at most the order of t , however by Lemma 7.2 and the fish Lemma 2.1 this separation then contradicts the choice of t .

Moreover, there exists such an efficient P - Q -distinguisher t which has an orientation \overrightarrow{t} such that $\overleftarrow{s} \leq \overrightarrow{t}$ for every $\overleftarrow{s} \in \overleftarrow{L}_i$: Otherwise $\overleftarrow{s} \leq \overrightarrow{t}$ for some orientation of t and if neither $P = P_s$ nor $Q = P_s$, then both P and Q would weakly orient t as \overleftarrow{t} since they weakly contain \overleftarrow{s} . On the other hand if $P = P_s$, say, then, again by (ii), the order of t is at least the order of s , thus s itself would be the required efficient P - Q -distinguisher.

Now consider, for every T_i , the set U^i of all separations t in U nested with \overleftarrow{L}_i and fulfilling the additional property of having, for every $\overleftarrow{s} \in \overleftarrow{L}_i$, an orientation such that $\overleftarrow{s} \leq \overrightarrow{t}$, i.e. U^i is the set of all separations in U inside of \overleftarrow{L}_i . \overrightarrow{U}^i is closed under \vee and \wedge in \overrightarrow{U} by the fish Lemma 2.1, thus the restriction of U to U^i is again a submodular universe of separations.

Given any $\overleftarrow{s} \in \overleftarrow{L}_i$, the down-closure of \overleftarrow{s} is a regular profile of U^i . Note that every efficient distinguisher for the profiles induced by \overleftarrow{s}_1 and $\overleftarrow{s}_2 \in \overleftarrow{L}_i$ on U^i is also an efficient distinguisher of P_{s_1} and P_{s_2} .

By Theorem 5.1 applied to the set of all separations of order less than k in U^i , we thus find a U^i -tree \hat{T}^i over \mathcal{F}_e (defined for \mathcal{P}_i). The corresponding nested set N_i efficiently distinguishes all these profiles induced by some $\overleftarrow{s}_i \in \overleftarrow{L}_i$.

But now the nested N given by $\bigcup_{i=1}^m (N_i \cup L_i)$ is as desired: It is easy to see that this set is nested and every N_i efficiently distinguishes any two profiles in \mathcal{P}_i . Moreover, we only ever changed separations inside of \overleftarrow{L}_i for every T_i .

The set N also contains an efficient P - Q -distinguisher for profiles P and Q in different T_i s: A profile R whose node in T lies on the path between the nodes containing P and Q , respectively, also does so in the tree induced by N . Thus, if we have efficient distinguishers for P and R and for R and Q , respectively, in N , then one of the two is also an efficient P - Q -distinguisher. An inductive application of this argument proves the claim, that the set N efficiently distinguishes any two profiles in \mathcal{P} . □

References

- [1] Diestel, R. (2017). *Graph Theory*, 5th ed. *Graduate Texts in Mathematics* **173**. Springer. [MR3644391](#)
- [2] Diestel, R. (2018). Abstract separation systems. *Order* **35** 157–170. [MR3774512](#)

- [3] Diestel, R., Eberenz, P. and Erde, J. (2017). Duality theorem for blocks and tangles in graphs. *SIAM J. Discrete Math.* **31** 1514–1528. [MR3670715](#)
- [4] Diestel, R., Erde, J. and Weißauer, D. (2019). Structural submodularity and tangles in abstract separation systems. *J. Combin. Theory (Series A)* **167C** 155–180. [MR3950083](#)
- [5] Diestel, R., Hundertmark, F. and Lemanczyk, S. (2019). Profiles of separations: in graphs, matroids, and beyond. *Combinatorica* **39** 37–75. [MR3936191](#)
- [6] Diestel, R. and Oum, S. (2019). Tangle-tree duality in graphs, matroids and beyond. *Combinatorica* **39** 879–910. [MR4015355](#)
- [7] Diestel, R. and Oum, S. (2021). Tangle-tree duality in abstract separation systems. *Advances in Mathematics* **377** 107470. [MR4186013](#)
- [8] Eberenz, P. Characteristics of profiles. MSc dissertation, Hamburg 2015.
- [9] Elbracht, C. and Kneip, J. A canonical tree-of-tangles theorem for submodular separation systems. Preprint, [arXiv:2009.02091](#).
- [10] Elbracht, C., Kneip, J. and Teegen, M. (2021). Trees of tangles in abstract separation systems. *J. Combin. Theory (Series A)* **180** 105425. [MR4213656](#)
- [11] Geelen, J., Gerards, B. and Whittle, G. (2009). Tangles, tree-decompositions and grids in matroids. *J. Combin. Theory (Series B)* **99** 657–667. [MR2518198](#)
- [12] Mazoit, F. (2013). A simple proof of the tree-width duality theorem. Preprint, [arXiv:1309.2266](#).
- [13] Robertson, N. and Seymour, P. D. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)* **52** 153–190. [MR1110468](#)

CHRISTIAN ELBRACHT
UNIVERSITÄT HAMBURG
FACHBEREICH MATHEMATIK
BUNDESSTRASSE 55
20146 HAMBURG
GERMANY
E-mail address: christian.elbracht@uni-hamburg.de

JAKOB KNEIP
UNIVERSITÄT HAMBURG
FACHBEREICH MATHEMATIK
BUNDESSTRASSE 55
20146 HAMBURG
GERMANY
E-mail address: jakob.kneip@uni-hamburg.de

MAXIMILIAN TEEGEN
UNIVERSITÄT HAMBURG
FACHBEREICH MATHEMATIK
BUNDESSTRASSE 55
20146 HAMBURG
GERMANY
E-mail address: maximilian.teegen@uni-hamburg.de

RECEIVED FEBRUARY 23, 2021