Structural submodularity and tangles in abstract separation systems

Reinhard Diestel, Joshua Erde and Daniel Weißauer

Abstract

We prove a tree-of-tangles theorem and a tangle-tree duality theorem for abstract separation systems $\vec{S}$ that are submodular in the structural sense that, for every pair of oriented separations, $\vec{S}$ contains either their meet or their join defined in some universe $\vec{U}$ of separations containing $\vec{S}$.

This holds, and is widely used, if $\vec{U}$ comes with a submodular order function and $\vec{S}$ consists of all its separations up to some fixed order. Our result is that for the proofs of these two theorems, which are central to abstract tangle theory, it suffices to assume the above structural consequence for $\vec{S}$, and no order function is needed.

1 Introduction

This paper is, in a sense, the capstone of a comprehensive project [2, 3, 4, 5, 8, 9, 10, 11, 13, 12, 17, 18] whose aim has been to utilize the idea of tangles familiar from Robertson and Seymour's graph minors project as a way of capturing clusters in other contexts, such as image analysis [14], genetics [6], or the social sciences [7]. The idea is to use tangles, which in graphs are certain consistent ways of orienting their low-order separations, as an indirect way of capturing 'fuzzy' clusters – ones that cannot easily be described by simply listing their elements – by instead orienting all those low-order separations towards them. We can then think of these as a collection of signposts all pointing to that cluster, and of clusters as collective targets of such consistent pointers.

Once clusters have been captured by 'abstract tangles' in this way, one can hope to generalize to such clusters Robertson and Seymour's two fundamental results about tangles in graphs [21]. One of these is the *tree-of-tangles* theorem. It says that any set of distinguishable tangles – ones that pairwise do not contain each other – can in fact be distinguished pairwise by a small and nested set of separations: for every pair of tangles there is a separation in this small and nested collection that distinguishes them. Formally, this means that these two tangles orient it differently; informally it means that one of its two orientations points to one of the tangles, while its other orientation points to the other tangle. Since these separations are nested, they split the underlying structure in a tree-like way, giving it a rough overall structure.

The other fundamental result from [21], the *tangle-tree duality* theorem, tells us that if there are no tangles of a desired type then the entire underlying structure can be split in such a tree-like way, i.e. by some nested set of separations, so that the regions corresponding to a node of the structure tree are all small. (What exactly this means may depend on the type of tangle considered.)

This research programme required a number of steps, of which this paper constitutes the last.

The first step was to make the notion of tangles independent from their natural habitat of graphs. In a graph, tangles are ways of consistently orienting all its separations $\{A, B\}$ up to some given order, either as (A, B) or as (B, A). If we want to do this for another kind of underlying structure than a graph, this structure will have to come with a notion of 'separation', it must be possible to 'orient' these separations, and there must be a difference between doing this 'consistently' or 'inconsistently'. If we wish to express, and perhaps prove, the two fundamental tangle theorems in such an abstract context, we further need a notion of when two 'separations' are nested.

There are many structures that come with a natural notion of separation. For sets, for example, we might simply take bipartitions. The notion of nestedness can then be borrowed from the nestedness of sets and applied to the bipartition classes. Thinking of a bipartition as an unordered pair of subsets, we can also naturally orient it 'towards one or the other of these subsets' by ordering the pair. Finally, we have to come up with natural notions of when orientations of different separations are consistent: we think of this as 'roughly pointing the same way', and it is another prerequisite for defining tangles to make this formal. This is both trickier to do in an abstract context and one of our main sources of freedom; we shall address this question in Section 2.

The completion of the first step in our research programme thus consisted in abstracting from the various notions of separation, and of consistently orienting separations, a minimum set of requirements that might serve as axioms for an abstract notion of tangle applicable to all of them. This resulted in the concept of *separation systems* and their ('abstract') tangles [8].

The second step, then, was to generalize the proofs of the tree-of-tangles theorem and the tangle-tree duality theorem to the abstract setting of separation systems. This was done in [11] and [13], respectively.

In order to prove these theorems, or to apply them to concrete cases of abstract separation systems, e.g. as in [12, 14], one so far still needed a further ingredient of graph tangles: a submodular order function on the separation system considered. Our aim in this paper is to show that one can do without this: we shall prove that a structural consequence of the existence of a submodular order function, a consequence that can be expressed in terms of abstract separation systems, can replace the assumption that such a function exists in the proofs of the above two theorems. We shall refer to separation systems that satisfy this structural condition as *submodular separation systems*.[1]

[1] There is also a notion of submodularity for separation *universes*. Separation universes are special separation systems that are particularly large, and they are always submodular

With this third step, then, the programme sketched above will be complete: we shall have a notion of tangle for very general abstract separation systems, as well as a tree-of-tangles theorem and a tangle-tree duality theorem for these tangles that can be expressed and proved without the need for any submodulary order function on the separation systems considered.

Formally, our two main results read as follows:

Theorem 1. *Every submodular separation system $\vec{S}$ contains a tree set of separations that distinguishes all the abstract tangles of S.*

Theorem 2. *Let $\vec{S}$ be a submodular separation system without degenerate elements in a distributive universe $\vec{U}$. Then exactly one of the following holds:*

(i) *S has an abstract tangle.*

(ii) *There exists an S-tree over $\mathcal{T}^*$ (witnessing that S has no abstract tangle).*

(See Section 2 for definitions.) Three further theorems, which partly strengthen or generalize the above two, will be stated in Section 2 (and proved later) when we have more terminology available.

One may ask, of course, whether weakening the existence of a submodular order function to 'structural submodularity' in the premise of these two theorems is worth the effort. We believe it is. For a start, the entire programme of developing abstract separation systems, and a theory of tangles for them, served the purpose of identifying the few structural assumptions one has to make of a set of objects called 'separations' in order to capture the essence of tangles in graphs, and thereby make them applicable in much wider contexts. It would then seem oblivious of these aims to stop just short of the goal: to continue to make unnecessarily strong assumptions of an extraneous and non-structural kind when weaker structural assumptions can achieve the same.

However, there is also a technical advantange. As we shall see in Sections 5.2 and 5.4, there are interesting abstract separation systems that are structurally submodular but which do not come with a natural submodular order function that implies this.

2 Abstract separation systems

Abstract separation systems were first introduced in [8]; see there for a gentle formal introduction and any terminology we forgot to define below. Motivation for why they are interesting can be found in the introductory sections of [11, 13, 12] and in [14]. In what follows we provide a self-contained account of just the definitions and basic facts about abstract separation systems that we need in this paper.

A *separation system* $(\vec{S}, \le, {}^*)$ is a partially ordered set with an order-reversing involution ${}^*\colon \vec{S} \to \vec{S}$. The elements of $\vec{S}$ are called *(oriented) separations*. The *inverse* of $\vec{s} \in \vec{S}$ is $\vec{s}^{\,*}$, which we usually denote by $\overleftarrow{s}$. An *(unoriented) sep-*

as separation systems. For separation universes, therefore, submodularity is used with the narrower meaning of being endowed with a submodular order function [8].

aration is a set $s = \{\vec{s}, \overleftarrow{s}\}$ consisting of a separation and its inverse and we then refer to $\vec{s}$ and $\overleftarrow{s}$ as the two *orientations* of s. Note that it may occur that $\vec{s} = \overleftarrow{s}$, we then call $\vec{s}$ *degenerate.* The set of all separations is denoted by S. When the context is clear, we often refer to oriented separations simply as separations in order to improve the flow of text.

If the partial order $(\vec{S}, \leq)$ is a lattice with join $\vee$ and meet $\wedge$, then we call $(\vec{S}, \leq, ^*, \vee, \wedge)$ a *universe* of (oriented) separations. It is *distributive* if it is distributive as a lattice. Typically, the separation systems we are interested in are contained in a universe of separations. In most applications, one starts with a universe $(\vec{U}, \leq, ^*, \vee, \wedge)$ and then defines $\vec{S}$ as a set of separations of low order with respect to some *order function* on $\vec{U}$, a map $|\cdot|\colon \vec{U} \to [0, \infty)$ that is *symmetric* in that $|\vec{s}| = |\overleftarrow{s}|$, and *submodular* in that $|\vec{s} \vee \vec{t}| + |\vec{s} \wedge \vec{t}| \leq |\vec{s}| + |\vec{t}|$ for all $\vec{s}, \vec{t} \in \vec{U}$. Submodularity of the order function in fact plays a crucial role in several arguments. One of its most immediate consequences is that whenever both $\vec{s}, \vec{t} \in \vec{S_k} := \{\vec{u} \in \vec{U} \colon |\vec{u}| < k\}$, then at least one of $\vec{s} \vee \vec{t}$ and $\vec{s} \wedge \vec{t}$ again lies in $\vec{S_k}$.

In order to avoid recourse to the external concept of an order function if possible, let us turn this last property into a definition that uses only the language of lattices. Let us call a subset M of a lattice $(L, \vee, \wedge)$ *submodular* if for all $x, y \in M$ at least one of $x \vee y$ and $x \wedge y$ lies in M. A separation system $\vec{S}$ contained in a given universe $\vec{U}$ of separations is *(structurally) submodular* if it is submodular as a subset of the lattice underlying $\vec{U}$.

We say that $\vec{s} \in \vec{S}$ is *small* (and $\overleftarrow{s}$ is co-small) if $\vec{s} \leq \overleftarrow{s}$. An element $\vec{s} \in \vec{S}$ is *trivial* in $\vec{S}$ (and $\overleftarrow{s}$ is co-trivial) if there exists $t \in S$ whose orientations $\vec{t}, \overleftarrow{t}$ satisfy $\vec{s} < \vec{t}$ as well as $\vec{s} < \overleftarrow{t}$. Notice that trivial separations are small.

Two separations $s, t \in S$ are *nested* if there exist orientations $\vec{s}$ of s and $\vec{t}$ of t such that $\vec{s} \leq \vec{t}$. Two oriented separations are nested if their underlying separations are. We say that two separations *cross* if they are not nested. A set of (oriented) separations is *nested* if any two of its elements are. A nested separation system without trivial or degenerate elements is a *tree set.* A set σ of non-degenerate oriented separations is a *star* if for any two distinct $\vec{s}, \vec{t} \in \sigma$ we have $\vec{s} \leq \overleftarrow{t}$. A family $\mathcal{F} \subseteq 2^{\vec{U}}$ of sets of separations is *standard for* $\vec{S}$ if for any trivial $\vec{s} \in \vec{S}$ we have $\{\overleftarrow{s}\} \in \mathcal{F}$. Given $\mathcal{F} \subseteq 2^{\vec{U}}$, we write $\mathcal{F}^*$ for the set of all elements of $\mathcal{F}$ that are stars.

An *orientation* of S is a set $O \subseteq \vec{S}$ which contains for every $s \in S$, exactly one of $\overleftarrow{s}, \vec{s}$. An orientation O of S is *consistent* if whenever $r, s \in S$ are distinct and $\vec{r} \leq \vec{s} \in O$, then $\overleftarrow{r} \notin O$. The idea behind this is that separations $\overleftarrow{r}$ and $\vec{s}$ are thought of as pointing away from each other if $\vec{r} \leq \vec{s}$. If we wish to orient r and s towards some common region of the structure which they are assumed to 'separate', as is the idea behind tangles, we should therefore not orient them as $\overleftarrow{r}$ and $\vec{s}$.

Tangles in graphs also satisfy another, more subtle, consistency requirement: they never orient three separations r, s, t so that the region to which they point collectively is 'small'.[2] This can be mimicked in abstract separation systems by

[2]Formally: so that the union of their sides to which they do *not* point is the entire graph.

asking that three oriented separations in an 'abstract tangle' must never have a co-small supremum; see [8, Section 5]. So let us implement this formally.

Given a family $\mathcal{F} \subseteq 2^{\vec{U}}$, we say that O *avoids* $\mathcal{F}$ if there is no $\sigma \subseteq O$ with $\sigma \in \mathcal{F}$. A consistent $\mathcal{F}$-avoiding orientation of S is called an *$\mathcal{F}$-tangle* of S. An $\mathcal{F}$-tangle for $\mathcal{F} = \mathcal{T}$ with

$$\mathcal{T} := \{\{\vec{r}, \vec{s}, \vec{t}\} \subseteq \vec{U} \colon \vec{r} \vee \vec{s} \vee \vec{t} \text{ is co-small}\}$$

is an *abstract tangle.*

A separation $s \in S$ *distinguishes* two orientations O_1, O_2 of S if $O_1 \cap s \neq O_2 \cap s$. Likewise, a set N of separations *distinguishes* a set $\mathcal{O}$ of orientations if for any two $O_1, O_2 \in \mathcal{O}$, there is some $s \in N$ which distinguishes them.

Let us restate our tree-of-tangles theorem for abstract tangles of submodular separation systems:

Theorem 1. *Every submodular separation system $\vec{S}$ contains a tree set of separations that distinguishes all the abstract tangles of S.*

We now introduce the structural dual to the existence of abstract tangles. An *S-tree* is a pair (T, α) consisting of a tree T and a map $\alpha : \vec{E}(T) \to \vec{S}$ from the set $\vec{E}(T)$ of orientations of edges of T to $\vec{S}$ such that $\alpha(y, x) = \alpha(x, y)^*$ for all $xy \in E(T)$. Given $\mathcal{F} \subseteq 2^{\vec{U}}$, we call (T, α) an *S-tree over $\mathcal{F}$* if $\alpha(F_t) \in \mathcal{F}$ for every $t \in T$, where

$$F_t := \{(s, t) \colon st \in E(T)\}.$$

It is easy to see that if S has an abstract tangle, then there can be no S-tree over $\mathcal{T}$.

Our tangle-tree duality theorem for abstract tangles of submodular separation systems, which we now re-state, asserts a converse to this. Recall that $\mathcal{T}^*$ denotes the set of stars in $\mathcal{T}$.

Theorem 2. *Let $\vec{S}$ be a submodular separation system without degenerate elements in a distributive universe $\vec{U}$. Then exactly one of the following holds:*

(i) *S has an abstract tangle.*

(ii) *There exists an S-tree over $\mathcal{T}^*$.*

Here, it really is necessary to exclude degenerate separations: a single degenerate separation will make the existence of abstract tangles impossible, although there might still be $\mathcal{T}^*$-tangles (and therefore no S-trees over $\mathcal{T}^*$). We will actually prove a duality theorem for $\mathcal{T}^*$-tangles without this additional assumption and then observe that $\mathcal{T}^*$-tangles are in fact already abstract tangles, unless $\vec{S}$ contains a degenerate separation.

In applications, we do not always wish to consider all the abstract tangles of a given separation system. For example, if S consists of the bipartitions $\{A, B\}$ of some finite set V (see [8] for definitions), then every $v \in V$ induces an abstract tangle

$$\tau_v := \big\{(A, B) \in \vec{S} \colon v \in B\big\},$$

the *principal tangle* induced by v. In particular, abstract tangles trivially exist in these situations. In order to exclude principal tangles, we could require that every tangle τ of S must satisfy $(\{v\}, V \setminus \{v\}) \in \tau$ for every $v \in V$.

More generally, we might want to prescribe for some separations s of S that any tangle of S we consider must contain a particular one of the two orientations of s rather than the other. This can easily be done in our abstract setting, as follows. Given $Q \subseteq \vec{U}$, let us say that an abstract tangle τ of S *extends* Q if $Q \cap \vec{S} \subseteq \tau$. It is easy to see that τ extends Q if and only if τ is $\mathcal{F}_Q$-avoiding, where

$$\mathcal{F}_Q := \{\{\overleftarrow{s}\} \colon \vec{s} \in Q \text{ non-degenerate}\}.$$

We call $Q \subseteq \vec{U}$ *down-closed* if $\vec{r} \le \vec{s} \in Q$ implies $\vec{r} \in Q$ for all $\vec{r}, \vec{s} \in \vec{U}$.

Here, then, is our refined tangle-tree duality theorem for abstract tangles of submodular separation systems.

Theorem 3. *Let $\vec{S}$ be a submodular separation system without degenerate elements in a distributive universe $\vec{U}$ and let $Q \subseteq \vec{U}$ be down-closed. Then exactly one of the following assertions holds:*

(i) *S has an abstract tangle extending Q.*

(ii) *There exists an S-tree over $\mathcal{T}^* \cup \mathcal{F}_Q$.*

Observe that Theorem 3 implies Theorem 2 by taking $Q = \emptyset$.

The abstract tangles in Theorem 3 are not the only $\mathcal{F}$-tangles for which such a statement holds. In [12] the tangle-tree duality theorem of [13] is used to prove such a statement for a broad class of $\mathcal{F}$-tangles, albeit under a stronger assumption: one needs there that $\vec{S}$ is not just structurally submodular, as is our assumption here throughout our paper, but that $\vec{U}$ has a submodular order function and $\vec{S}$ is the set separations up to some fixed order (and therefore, in particular, submodular).

In Section 4, however, we will show that the weaker assumption that $\vec{S}$ itself is submodular is in fact sufficient to establish the only property of S whose proof in [12] requires a submodular order function: this is the fact that S is 'separable'. (We shall repeat the definition of this in Section 4.)

The other ingredient one needs for all those applications of the tangle-tree duality theorem from [13] is a property of $\mathcal{F}$: that $\mathcal{F}$ is 'closed under shifting'. Sometimes, a submodular order function on $\vec{U}$ is needed also to establish this property of $\mathcal{F}$. But if it is not, we can now prove the same application without a submodular order function, assuming only that S itself is submodular:

Theorem 4. *Let $\vec{U}$ be a universe of separations and $\vec{S} \subseteq \vec{U}$ a submodular separation system. Let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars which is standard for $\vec{S}$ and closed under shifting. Then exactly one of the following holds:*

(i) *There exists an $\mathcal{F}$-tangle of S.*

(ii) *There exists an S-tree over $\mathcal{F}$.*

We shall prove Theorem 4 in Section 4.

Our last result is an example of Theorem 4 for a concrete $\mathcal{F}$, a tangle-tree duality theorem for $\mathcal{F}$-tangles of bipartitions of a set that are used particularly often in applications [6, 7]. Let $\vec{U}$ be the universe of oriented bipartitions (A, B) of a set V (see [8] for definitions). Let $m \geq 1$ and $n \geq 2$ be integers, and define

$$\mathcal{F}_m := \{\, F \subseteq \vec{U} : \big| \bigcap_{(A,B)\in F} B \,\big| < m \,\}$$

and

$$\mathcal{F}_m^n := \{\, F \in \mathcal{F}_m : |F| < n \,\}.$$

To subsume $\mathcal{F}_m$ under this latter notation we allow $n = \infty$, so that $\mathcal{F}_m^\infty = \mathcal{F}_m$. Given any collection $\mathcal{F} \subseteq 2^{\vec{U}}$ of sets of oriented separations, we write $\mathcal{F}^*$ for its subcollection of those sets $F \in \mathcal{F}$ that are stars (of oriented separations).

We shall prove in Section 4 that the set of stars in $\mathcal{F}_m$ is closed under shifting. Building on Theorem 4, we then use this in Section 5 to prove the following:

Theorem 5. *Let $\vec{S} \subseteq \vec{U}$ be a submodular separation system, let $1 \leq m \in \mathbb{N}$ and $2 \leq n \in \mathbb{N} \cup \{\infty\}$, and let $\mathcal{F} = \mathcal{F}_m^n$. Then exactly one of the following two statements holds:*

(i) *S has an $\mathcal{F}$-tangle;*

(ii) *There exists an S-tree over $\mathcal{F}^*$.*

The bound n on the size of the sets in $\mathcal{F}$ is often taken to be 4. In (i) we could replace $\mathcal{F}$ with $\mathcal{F}^*$, since for these $\mathcal{F}$ the $\mathcal{F}$-tangles are precisely the $\mathcal{F}^*$-tangles; see Section 5.3.

3 The tree-of-tangles theorem

In this section we will prove Theorem 1. In fact, we are going to prove a slightly more general statement. Let $\mathcal{P} := \{\{\vec{s}, \vec{t}, (\vec{s} \vee \vec{t})^*\} : \vec{s}, \vec{t} \in \vec{U}\}$. The $\mathcal{P}$-tangles are known as *profiles*. A profile of S is *regular* if it contains all the small separations in $\vec{S}$.

Theorem 6. *Let $\vec{S}$ be a submodular separation system and Π a set of profiles of S. Then $\vec{S}$ contains a tree set that distinguishes Π.*

This implies Theorem 1, by the following easy observation.

Lemma 7. *Every abstract tangle is a profile.*

Proof. Let $\vec{s}, \vec{t} \in \vec{U}$ and $\vec{r} := \vec{s} \vee \vec{t}$. Then

$$\vec{s} \vee \vec{t} \vee \overleftarrow{r} = \vec{r} \vee \overleftarrow{r}$$

is co-small, so $\{\vec{s}, \vec{t}, \overleftarrow{r}\} \in \mathcal{T}$. Therefore $\mathcal{P} \subseteq \mathcal{T}$ and every $\mathcal{T}$-tangle is also a $\mathcal{P}$-tangle. $\square$

We first recall a basic fact about nestedness of separations. For $s, t \in S$, we define the *corners* $\vec{s} \wedge \vec{t}$, $\vec{s} \wedge \overleftarrow{t}$, $\overleftarrow{s} \wedge \vec{t}$ and $\overleftarrow{s} \wedge \overleftarrow{t}$.

Lemma 8 ([8])**.** *Let $\vec{S}$ be a separation system in a universe $\vec{U}$ of separations. Let s, t be two crossing separations and $\vec{r}$ one of the corners. Then every separation that is nested with both s and t is nested with r as well.*

In the proof of Theorem 6, we take a nested set $\mathcal{N}$ of separations that distinguishes some set Π_0 of regular profiles and we want to exchange one element of $\mathcal{N}$ by some other separation while maintaining that Π_0 is still distinguished. The following lemma simplifies this exchange.

Lemma 9. *Let $\vec{S}$ be a separation system, $\mathcal{O}$ a set of consistent orientations of S and $\mathcal{N} \subseteq S$ an inclusion-minimal nested set of separations that distinguishes $\mathcal{O}$. Then for every $t \in \mathcal{N}$ there is a unique pair of orientations $O_1, O_2 \in \mathcal{O}$ that are distinguished by t and by no other element of $\mathcal{N}$.*

Proof. It is clear that at least one such pair must exist, for otherwise $\mathcal{N} \setminus \{t\}$ would still distinguish $\mathcal{O}$, thus violating the minimality of $\mathcal{N}$.

Suppose there was another such pair, say O_1', O_2'. After relabeling, we may assume that $\vec{t} \in O_1 \cap O_1'$ and $\overleftarrow{t} \in O_2 \cap O_2'$. By symmetry, we may further assume that $O_1 \neq O_1'$. Since $\mathcal{N}$ distinguishes $\mathcal{O}$, there is some $r \in \mathcal{N}$ with $\vec{r} \in O_1$, $\overleftarrow{r} \in O_1'$.

As t is the only element of $\mathcal{N}$ distinguishing O_1, O_2, it must be that $\vec{r} \in O_2$ as well, and similarly $\overleftarrow{r} \in O_2'$. We hence see that for any orientation τ of $\{r, t\}$, there is an $O \in \{O_1, O_2, O_1', O_2'\}$ with $\tau \subseteq O$. Since $\mathcal{N}$ is nested, there exist orientations of r and t pointing away from each other. But then one of O_1, O_2, O_1', O_2' is inconsistent, which is a contradiction. $\square$

Proof of Theorem 6. Note that it suffices to show that there is a nested set $\mathcal{N}$ of separations that distinguishes Π: Every consistent orientation contains every trivial and every degenerate element, so any inclusion-minimal such set $\mathcal{N}$ gives rise to a tree-set.

We prove this by induction on $|\Pi|$, the case $|\Pi| = 1$ being trivial.

For the induction step, let $P \in \Pi$ be arbitrary and $\Pi_0 := \Pi \setminus \{P\}$. By the induction hypothesis, there exists a nested set $\mathcal{N}$ of separations that distinguishes Π_0. If some such set $\mathcal{N}$ distinguishes Π, there is nothing left to show. Otherwise, for every nested $\mathcal{N} \subseteq S$ which distinguishes Π_0 there is a $P' \in \Pi_0$ which $\mathcal{N}$ does not distinguish from P. Note that P' is unique. For any $s \in S$ that distinguishes P and P', let $d(\mathcal{N}, s)$ be the number of elements of $\mathcal{N}$ which are not nested with s.

Choose a pair $(\mathcal{N}, s)$ so that $d(\mathcal{N}, s)$ is minimum. Clearly, we may assume $\mathcal{N}$ to be inclusion-minimal with the property of distinguishing Π_0. If $d(\mathcal{N}, s) = 0$, then $\mathcal{N} \cup \{s\}$ is a nested set distinguishing Π and we are done, so we now assume for a contradiction that $d(\mathcal{N}, s) > 0$.

Since $\mathcal{N}$ does not distinguish P and P', we can fix an orientation of each $t \in \mathcal{N}$ such that $\vec{t} \in P \cap P'$. Choose a $t \in \mathcal{N}$ such that t and s cross and $\vec{t}$ is minimal. Let (P_1, P_2) be the unique pair of profiles in Π_0 which are distinguished by t and by no other element of $\mathcal{N}$, say $\overleftarrow{t} \in P_1$, $\vec{t} \in P_2$. Let us assume without loss of generality that $\overleftarrow{s} \in P_1$. The situation is depicted in Figure 1. Note that we do not know whether $\vec{s} \in P_2$ or $\overleftarrow{s} \in P_2$. Also, the roles of P and P' might be reversed, but this is insignificant.

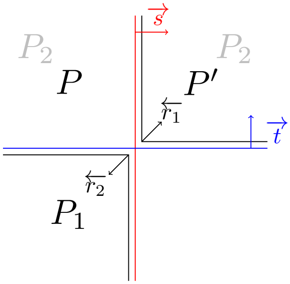

Figure 1: Crossing separations

Suppose first that $\vec{r_1} := \vec{s} \vee \vec{t} \in \vec{S}$. Let $Q \in \{P, P'\}$. If $\vec{s} \in Q$, then $\vec{r_1} \in Q$, since $\vec{t} \in P \cap P'$ and Q is a profile. If $\vec{r_1} \in Q$, then $\vec{s} \in Q$ since Q is consistent and $\vec{s} \le \vec{r_1} \in Q$: it cannot be that $\vec{s} = \overleftarrow{r_1}$, since then s and t would be nested. Hence each $Q \in \{P, P'\}$ contains $\vec{r_1}$ if and only if it contains $\vec{s}$. In particular, r_1 distinguishes P and P'. By Lemma 8, every $u \in \mathcal{N}$ that is nested with s is also nested with r_1. Moreover, t is nested with r_1, but not with s, so that $d(\mathcal{N}, r_1) < d(\mathcal{N}, s)$. This contradicts our choice of s.

Therefore $\vec{s} \vee \vec{t} \notin \vec{S}$. Since $\vec{S}$ is submodular, it follows that $\vec{r_2} := \vec{s} \wedge \vec{t} \in \vec{S}$. Moreover, r_2 is nested with every $u \in \mathcal{N} \setminus \{t\}$. This is clear if $\vec{t} \le \vec{u}$ or $\vec{t} \le \overleftarrow{u}$, since $\vec{r_2} \le \vec{t}$. It cannot be that $\overleftarrow{u} \le \vec{t}$, because $\vec{u}, \vec{t} \in P$ and P is consistent. Since $\mathcal{N}$ is nested, only the case $\vec{u} < \vec{t}$ remains. Then, by our choice of $\vec{t}$, u and s are nested and it follows from Lemma 8 that u and r_2 are also nested. Hence $\mathcal{N}' := (\mathcal{N} \setminus \{t\}) \cup \{r_2\}$ is a nested set of separations.

To see that $\mathcal{N}'$ distinguishes Π_0, it suffices to check that r_2 distinguishes P_1 and P_2. We have $\vec{r_2} \in P_2$ since P_2 is consistent and $\vec{r_2} \le \vec{t} \in P_2$: if $\vec{r_2} = \overleftarrow{t}$, then s and t would be nested. Since $\overleftarrow{r_2} = \overleftarrow{s} \vee \overleftarrow{t}$ and $\overleftarrow{s}, \overleftarrow{t} \in P_1$, we find $\overleftarrow{r_2} \in P_1$. Any element of $\mathcal{N}'$ which is not nested with s lies in $\mathcal{N}$. Since $t \in \mathcal{N} \setminus \mathcal{N}'$ is not nested with s, it follows that $d(\mathcal{N}', s) < d(\mathcal{N}, s)$, contrary to our choice of $\mathcal{N}$ and s. $\square$

4 Tangle-tree duality

Our agenda for this section is first to prove Theorem 4, and then to derive from it Theorem 3, which as we have seen implies Theorem 2. Our proof will be an application of the basic tangle-tree duality theorem from [13].

For this we need to introduce the notion of separability, and then prove that submodular separation systems are separable (Lemma 13). This lemma not only lies at the heart of our proof of Theorem 4: it will also be central to any other result that asserts a tangle-tree type duality for separation systems

$(\vec{S}, \le, ^*)$ that are structurally submodular, but are not so simply as a corollary of the existence of a submodular order function on $\vec{S}$.

A separation $\vec{s} \in \vec{S}$ *emulates* $\vec{r}$ *in* $\vec{S}$ if $\vec{s} \ge \vec{r}$ and for every $\vec{t} \in \vec{S} \setminus \{\overleftarrow{r}\}$ with $\vec{t} \ge \vec{r}$ we have $\vec{s} \vee \vec{t} \in \vec{S}$. For $\vec{s} \in \vec{S}$, $\sigma \subseteq \vec{S}$ and $\vec{x} \in \sigma$, define

$$\sigma^{\vec{s}}_{\vec{x}} := \{\vec{x} \vee \vec{s}\} \cup \{\vec{y} \wedge \overleftarrow{s} \colon \vec{y} \in \sigma \setminus \{\vec{x}\}\}.$$

Lemma 10. *Suppose $\vec{s} \in \vec{S}$ emulates a non-trivial $\vec{r}$ in $\vec{S}$, and let $\sigma \subseteq \vec{S}$ be a star such that $\vec{r} \le \vec{x} \in \sigma$. Then $\sigma^{\vec{s}}_{\vec{x}} \subseteq \vec{S}$ is a star.*

Proof. Note that for every $\vec{y} \in \sigma \setminus \{\vec{x}\}$ we have $\vec{r} \le \overleftarrow{y}$. It is clear that for any two distinct $\vec{u}, \vec{v} \in \sigma^{\vec{s}}_{\vec{x}}$ we have $\vec{u} \le \overleftarrow{v}$, so we only need to show that every element of $\sigma^{\vec{s}}_{\vec{x}}$ is non-degenerate and lies in $\vec{S}$. For every $\vec{u} \in \sigma^{\vec{s}}_{\vec{x}}$ there is a non-degenerate $\vec{t} \in \vec{S}$ with $\vec{r} \le \vec{t}$ such that either $\vec{u} = \vec{t} \vee \vec{s}$ or $\overleftarrow{u} = \vec{t} \vee \vec{s}$.

Let $\vec{t} \in \vec{S}$ be non-degenerate with $\vec{r} \le \vec{t}$. Since $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$, we find $\vec{t} \vee \vec{s} \in \vec{S}$. Assume for a contradiction that $\vec{t} \vee \vec{s}$ was degenerate. Since $\vec{t}$ is non-degenerate, we find that $\vec{t} < \vec{t} \vee \vec{s}$, so that $\vec{t}$ is trivial. But then so is $\vec{r}$, because $\vec{r} \le \vec{t}$. This contradicts our assumption on $\vec{r}$. □

The separation system $\vec{S}$ is *separable* if for all non-trivial and non-degenerate $\vec{r_1}, \overleftarrow{r_2} \in \vec{S}$ with $\vec{r_1} \le \vec{r_2}$ there exists an $\vec{s} \in \vec{S}$ which emulates $\vec{r_1}$ in $\vec{S}$ while simultaneously $\overleftarrow{s}$ emulates $\overleftarrow{r_2}$ in $\vec{S}$.

Given some $\mathcal{F} \subseteq 2^{\vec{U}}$, we say that $\vec{s}$ *emulates* $\vec{r}$ *in* $\vec{S}$ *for* $\mathcal{F}$ if $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$ and for every star $\sigma \subseteq \vec{S} \setminus \{\overleftarrow{r}\}$ with $\sigma \in \mathcal{F}$ and every $\vec{x} \in \sigma$ with $\vec{x} \ge \vec{r}$ we have $\sigma^{\vec{s}}_{\vec{x}} \in \mathcal{F}$.

The separation system $\vec{S}$ is $\mathcal{F}$*-separable* if for all non-trivial and non-degenerate $\vec{r_1}, \overleftarrow{r_2} \in \vec{S}$ with $\vec{r_1} \le \vec{r_2}$ and $\{\overleftarrow{r_1}\}, \{\vec{r_2}\} \notin \mathcal{F}$ there exists an $\vec{s} \in \vec{S}$ which emulates $\vec{r_1}$ in $\vec{S}$ for $\mathcal{F}$ while simultaneously $\overleftarrow{s}$ emulates $\overleftarrow{r_2}$ in $\vec{S}$ for $\mathcal{F}$.

Theorem 11 ([13, Theorem 4.3])**.** *Let $\vec{U}$ be a universe of separations and $\vec{S} \subseteq \vec{U}$ a separation system. Let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars, standard for $\vec{S}$. If $\vec{S}$ is $\mathcal{F}$-separable, then exactly one of the following holds:*

(i) *There exists an $\mathcal{F}$-tangle of S.*

(ii) *There exists an S-tree over $\mathcal{F}$.*

In applications of Theorem 11 it is often easier to split the proof of the main premise, that $\vec{S}$ is $\mathcal{F}$-separable, into two parts: a proof that $\vec{S}$ is separable and one that $\mathcal{F}$ is *closed under shifting* in $\vec{S}$: that whenever $\vec{s} \in \vec{S}$ emulates (in $\vec{S}$) some nontrivial and nondegenerate $\vec{r} \le \vec{s}$ not forced by $\mathcal{F}$, then it does so for $\mathcal{F}$. Indeed, the following is immediate from the definitions:

Lemma 12. *Let $\vec{U}$ be a universe of separations, $\vec{S} \subseteq \vec{U}$ a separation system, and $\mathcal{F} \subseteq 2^{\vec{U}}$ a set of stars. If $\vec{S}$ is separable and $\mathcal{F}$ is closed under shifting, then $\vec{S}$ is $\mathcal{F}$-separable.* □

It is shown in [12] that if $\vec{U}$ is a universe of separations with an order function, then the sets $\vec{S_k}$ of all separations of order less than some fixed positive integer k

are separable for all k, and virtually all the applications of Theorem 11 that are given in [13] involve a separation system of the form $\vec{S_k}$.

While many applications of the submodularity of an order function use only its structural consequence that motivated our abstract notion of submodularity, the use of submodularity in the proof that $\vec{S_k}$ is separable – see [12, Lemma 3.4] – uses it in a more subtle way. There, the orders of opposite corners of two crossing separations $\vec{s}$ and $\vec{t}$ are compared not with any fixed value of k but with the (possibly distinct) orders of s and t directly. This kind of argument is naturally difficult, if not impossible, to mimic in our set-up.

However, we can prove this nevertheless, choosing a different route. The following lemma is, in essence, the main result of this section:

Lemma 13. *Let $\vec{U}$ be a universe of separations and $\vec{S} \subseteq \vec{U}$ a submodular separation system. Then $\vec{S}$ is separable.*

We will actually prove a slightly more general statement about submodular lattices. Let $(L, \vee, \wedge)$ be a lattice and let $M \subseteq L$. Given $x, y \in M$, we say that x *pushes* y if $x \leq y$ and for any $z \in M$ with $z \leq y$ we have $x \wedge z \in M$. Similarly, we say that x *lifts* y if $x \geq y$ and for any $z \in M$ with $z \geq y$ we have $x \vee z \in M$. Observe that both of these relations are reflexive and transitive: Every $x \in M$ pushes (lifts) itself and if x pushes (lifts) y and y pushes (lifts) z, then x pushes (lifts) z. We say that M is *strongly separable* if for all $x, y \in M$ with $x \leq y$ there exists a $z \in M$ that lifts x and pushes y.

The definitions of lifting, pushing and separable extend verbatim to a separation system within a universe of separations when regarded as a subset of the underlying lattice. These notions are strengthenings of the notions of emulating and separable: If $\vec{s} \in \vec{S}$ lifts $\vec{r} \in \vec{S}$, then $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$, and $\vec{s}$ pushes $\vec{r}$ if and only if $\overleftarrow{s}$ lifts $\overleftarrow{r}$. Similarly, if $\vec{S}$ is strongly separable, then $\vec{S}$ is separable. Lemma 13 is then an immediate consequence of the following:

Lemma 14. *Let L be a finite lattice and $M \subseteq L$ submodular. Then M is strongly separable.*

Proof. Call a pair $(a, b) \in M \times M$ *bad* if $a \leq b$ and there is no $x \in M$ that lifts a and pushes b. Assume for a contradiction that there was a bad pair and choose one, say (a, b), such that $I(a, b) := \{u \in M : a \leq u \leq b\}$ is minimal.

We claim that a pushes every $z \in I(a, b) \setminus \{b\}$. Indeed, assume for a contradiction a did not push some such z. By minimality of (a, b), the pair (a, z) is not bad, so there is some $x \in M$ which lifts a and pushes z. By assumption, $x \neq a$ and so by minimality, the pair (x, b) is not bad, yielding a $y \in M$ which lifts x and pushes b. By transitivity, it follows that y lifts a. But then (a, b) is not a bad pair, which is a contradiction. An analogous argument establishes that b lifts every $z \in I(a, b) \setminus \{a\}$.

Since (a, b) is bad, a does not push b, so there is some $x \in M$ with $x \leq b$ for which $a \wedge x \notin M$. Similarly, there is a $y \in M$ with $y \geq a$ for which $b \vee y \notin M$. Since M is submodular, it follows that $a \vee x, b \wedge y \in M$. Note that $a \vee x, b \wedge y \in I(a, b)$. Furthermore, $x \leq a \vee x$ and $a \wedge x \notin M$, so a does not push $a \vee x$. We showed that a pushes every $z \in I(a,b) \setminus \{b\}$, so it follows that $a \vee x = b$. Similarly, we find that $b \wedge y = a$. But then

$$x \vee y = x \vee (a \vee y) = b \vee y \notin M,$$
$$x \wedge y = (x \wedge b) \wedge y = x \wedge a \notin M.$$

This contradicts the submodularity of M. □

As a result we obtain our tangle-tree duality theorem for $\mathcal{F}$-tangles of submodular separation systems, which we restate:

Theorem 4. *Let $\vec{U}$ be a universe of separations and $\vec{S} \subseteq \vec{U}$ a submodular separation system. Let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars which is standard for $\vec{S}$ and closed under shifting. Then exactly one of the following holds:*

(i) *There exists an $\mathcal{F}$-tangle of S.*

(ii) *There exists an S-tree over $\mathcal{F}$.*

Proof. Since $\vec{S}$ is submodular, Lemma 13 implies that $\vec{S}$ is separable. Since $\mathcal{F}$ is closed under shifting, it follows from Lemma 12 that $\vec{S}$ is $\mathcal{F}$-separable. The result then follows by Theorem 11. □

We will now use Theorem 4 to prove Theorem 3, which in turn implies Theorem 2. Recall that we are considering a downclosed subset $Q \subseteq \vec{U}$ of a distributive universe of separations, and a submodular separation system $\vec{S}$ without degenerate elements in $\subseteq \vec{U}$, and we wish to prove a tangle-tree duality theorem for abstract tangles of $\vec{S}$ extending Q. Note that these are precisely the $(\mathcal{T} \cup \mathcal{F}_Q)$-tangles of $\vec{S}$. However, since the family $\mathcal{F}$ in Theorem 4 is assumed to be a set of stars, we cannot work directly with $\mathcal{T}$. Instead we will work with $\mathcal{T}^*$, the set of stars in $\mathcal{T}$. It will turn out that, since $\vec{S}$ has no degenerate elements, this will not change the set of $\mathcal{T}$-tangles (cf. Lemma 17). So, we will first show that we can apply Theorem 4 with $\mathcal{F} = \mathcal{T}_Q$, where $\mathcal{T}_Q := \mathcal{T}^* \cup \mathcal{F}_Q$, and then show that the $\mathcal{T}^*$-tangles are precisely the $\mathcal{T}$-tangles. Theorem 3 will then follow.

Let us first prove the following simple fact, which will be useful in a few different situations.

Lemma 15. *Let $\vec{U}$ be a distributive universe of separations. Let $\vec{u}, \vec{v}, \vec{w} \in \vec{U}$. If $\vec{u} \le \vec{v}$ and $\vec{v} \vee \vec{w}$ is co-small, then $\vec{v} \vee (\vec{w} \wedge \overleftarrow{u})$ is co-small.*

Proof. Let $\vec{x} := \vec{v} \vee (\vec{w} \wedge \overleftarrow{u})$. By distributivity of $\vec{U}$

$$\vec{x} = (\vec{v} \vee \vec{w}) \wedge (\vec{v} \vee \overleftarrow{u}) \ge (\vec{v} \vee \vec{w}) \wedge (\vec{u} \vee \overleftarrow{u}).$$

Let $\vec{s} := \vec{v} \vee \vec{w}$ and $\vec{t} := \vec{u} \vee \overleftarrow{u}$. Then $\overleftarrow{s} \le \vec{s}$ by assumption and $\overleftarrow{s} \le \overleftarrow{v} \le \vec{t}$. Further $\overleftarrow{t} \le \vec{u} \le \vec{t}$ and $\overleftarrow{t} \le \vec{u} \le \vec{v}$. Therefore

$$\overleftarrow{x} \le \overleftarrow{s} \vee \overleftarrow{t} \le \vec{s} \wedge \vec{t} \le \vec{x}.$$ □

In order to apply Theorem 4 with $\mathcal{F} = \mathcal{T}_Q$, we need to show that T_Q is closed under shifting.

Lemma 16. *If $Q \subseteq \vec{U}$ is down-closed and $\vec{U}$ is distributive, then $\mathcal{T}_Q$ is closed under shifting.*

Proof. Let $\vec{r} \in \vec{S}$ non-trivial and non-degenerate with $\{\overleftarrow{r}\} \notin \mathcal{F}$. Let $\vec{s} \in \vec{S}$ emulate $\vec{r}$ in $\vec{S}$, let $\mathcal{T}_Q \ni \sigma \subseteq \vec{S} \setminus \{\overleftarrow{r}\}$ and $\vec{r} \leq \vec{x} \in \sigma$. We have to show that $\sigma^{\vec{s}}_{\vec{x}} \in \mathcal{T}_Q$. From Lemma 10 we know that $\sigma^{\vec{s}}_{\vec{x}}$ is a star, so we only need to verify that $\sigma^{\vec{s}}_{\vec{x}} \in \mathcal{T}^* \cup \mathcal{F}_Q$.

Suppose first that $\sigma \in \mathcal{T}^*$. Let $\vec{w} := \bigvee(\sigma \setminus \{\vec{x}\})$. Applying Lemma 15 with $\vec{u} = \vec{s}$ and $\vec{v} = \vec{x} \vee \vec{s}$, we see that

$$\bigvee \sigma^{\vec{s}}_{\vec{x}} = (\vec{x} \vee \vec{s}) \vee (\vec{w} \wedge \overleftarrow{s})$$

is co-small. Since $\sigma^{\vec{s}}_{\vec{x}}$ has at most three elements, it follows that $\sigma^{\vec{s}}_{\vec{x}} \in \mathcal{T}$.

Suppose now that $\sigma \in \mathcal{F}_Q$. Then $\sigma = \{\vec{x}\}$ and $\overleftarrow{x} \in Q$. As Q is down-closed, we have $\overleftarrow{x} \wedge \overleftarrow{s} \in Q$. Since $\sigma^{\vec{s}}_{\vec{x}}$ is a star, $\overleftarrow{x} \wedge \overleftarrow{s}$ is non-degenerate and therefore

$$\sigma^{\vec{s}}_{\vec{x}} = \{\vec{x} \vee \vec{s}\} = \{(\overleftarrow{x} \wedge \overleftarrow{s})^*\} \in \mathcal{F}_Q.$$

$\square$

Lemma 17. *Let $\vec{U}$ be a distributive universe of separations and let $\vec{S} \subseteq \vec{U}$ be a submodular separation system without degenerate elements. Then the $\mathcal{T}^*$-tangles are precisely the abstract tangles.*

Proof. Since $\mathcal{T}^* \subseteq \mathcal{T}$, every abstract tangle is also a $\mathcal{T}^*$-tangle. We only need to show that, conversely, every $\mathcal{T}^*$-tangle in fact avoids $\mathcal{T}$.

For $\sigma \in \mathcal{T}$, let $d(\sigma)$ be the number of pairs $\vec{s}, \vec{t} \in \sigma$ which are not nested. Let O be a consistent orientation of S and suppose O was not an abstract tangle. Choose $\mathcal{T} \ni \sigma \subseteq O$ such that $d(\sigma)$ is minimum and, subject to this, σ is inclusion-minimal. We will show that σ is indeed a star, thus showing that O is not a $\mathcal{T}^*$-tangle.

If σ contained two comparable elements, say $\vec{s} \leq \vec{t}$, then $\sigma' := \sigma \setminus \{\vec{s}\}$ satisfies $\sigma' \in \mathcal{T}$, $\sigma' \subseteq O$ and $d(\sigma') \leq d(\sigma)$, violating the fact that σ is inclusion-minimal. Hence σ is an antichain. Since $\vec{S}$ has no degenerate elements, it follows from the consistency of O that any two nested $\vec{s}, \vec{t} \in \sigma$ satisfy $\vec{s} \leq \overleftarrow{t}$. To show that σ is a star, it thus suffices to prove that any two elements are nested.

Suppose that σ contained two crossing separations, say $\vec{s}, \vec{t} \in \sigma$. By submodularity of $\vec{S}$, at least one of $\vec{s} \wedge \overleftarrow{t}$ and $\overleftarrow{s} \wedge \vec{t}$ lies in $\vec{S}$. By symmetry we may assume that $\vec{r} := \vec{s} \wedge \overleftarrow{t} \in \vec{S}$. Let $\sigma' := (\sigma \setminus \{\vec{s}\}) \cup \{\vec{r}\}$. Since O is consistent, $\vec{r} \leq \vec{s}$ and $r \neq s$, it follows that $\vec{r} \in O$ and so $\sigma' \subseteq O$ as well. Let $\vec{w} = \bigvee(\sigma \setminus \{\vec{t}\})$. As $\vec{t} \vee \vec{w} = \bigvee \sigma$ is co-small, we can apply Lemma 15 with $\vec{u} = \vec{v} = \vec{t}$ to deduce that $\vec{t} \vee (\vec{w} \wedge \overleftarrow{t})$ is co-small as well. But

$$\vec{t} \vee (\vec{w} \wedge \overleftarrow{t}) = \vec{t} \vee \bigvee_{\vec{x} \in \sigma \setminus \{\vec{t}\}} (\vec{x} \wedge \overleftarrow{t}) \leq \bigvee \sigma',$$

so $\bigvee \sigma'$ is also co-small and $\sigma' \in \mathcal{T}$.

We now show that $d(\sigma') < d(\sigma)$. Since s and t cross, while r and t do not, it suffices to show that every $\vec{x} \in \sigma \setminus \{\vec{s}\}$ which is nested with $\vec{s}$ is also nested with $\vec{r}$. But for every such $\vec{x}$ we have $\vec{s} \leq \overleftarrow{x}$. Since $\vec{r} \leq \vec{s}$, we get $\vec{r} \leq \overleftarrow{x}$ as well, showing that r and x are nested. So in fact $d(\sigma') < d(\sigma)$, which is a contradiction. This completes the proof that σ is nested and therefore a star. □

We are now in a position to prove Theorem 3.

Proof of Theorem 3. By Lemma 13, $\vec{S}$ is separable, and by Lemma 16, $\mathcal{T}_Q$ is closed under shifting. Therefore, by Theorem 4, there is no S-tree over $\mathcal{T}^* \cup \mathcal{F}_Q$, if and only if S has a T_Q-tangle, that is, a $\mathcal{T}^*$-tangle extending Q.

However, since $\vec{U}$ is distributive and $\vec{S}$ contains no degenerate elements, Lemma 17 implies that S has a $\mathcal{T}^*$-tangle extending Q if and only if S has an abstract tangle extending Q. □

5 Special cases and applications

5.1 Tangles in graphs and matroids

We briefly indicate how tangles in graphs and matroids can be seen as special cases of abstract tangles in separation systems. Tangles in graphs and hypergraphs were introduced by Robertson and Seymour in [21], but a good deal of the work is done in the setting of connectivity systems. Geelen, Gerards and Whittle [20] made this more explicit and defined tangles as well as the dual notion of branch-decompositions for connectivity systems, an approach that we will follow.

Let X be a finite set and $\lambda : 2^X \to \mathbb{Z}$ a map assigning integers to the subsets of X such that $\lambda(X \setminus A) = \lambda(A)$ for all $A \subseteq X$ and

$$\lambda(A \cup B) + \lambda(A \cap B) \leq \lambda(A) + \lambda(B)$$

for all $A, B \subseteq X$. The pair (X, λ) is then called a *connectivity system*.

Both graphs and matroids give rise to connectivity systems. For a given graph G, we can take $X := E(G)$ and define $\lambda(F)$ as the number of vertices of G incident with edges in both F and $E \setminus F$. Given a matroid M with ground-set X and rank-function r, we take λ to be the connectivity function $\lambda(A) := r(A) + r(X \setminus A) - r(X)$.

Now consider 2^X as a universe of separations with set-inclusion as the partial order and $A^* = X \setminus A$ as involution. For an integer k, the set $\vec{S_k}$ of all sets A with $\lambda(A) < k$ is then a submodular separation system. Let $Q := \{\emptyset\} \cup \{\{x\} : x \in X\}$ consist of the empty-set and all singletons of X and note that Q is down-closed.

A *tangle of order k* of (X, λ), as defined in [20], is then precisely an abstract tangle extending Q. It is easy to see that (X, λ) has a *branch-decomposition* of width $< k$ if and only if there exists an S_k-tree over $\mathcal{T}^* \cup \mathcal{F}_Q$. Theorem 3 then yields the classic duality theorem for tangles and branch-decompositions in connectivity systems, see [21, 20].

5.2 Clique separations

We now describe a submodular separation system that is not derived from a submodular order function, and provide a natural set of stars for which Theorem 4 applies.

Let $G = (V, E)$ be a finite graph and $\vec{U}$ the universe of all separations of G, that is, pairs (A, B) of subsets of V with $V = A \cup B$ such that there is no edge between $A \setminus B$ and $B \setminus A$. Here the partial order is given by $(A, B) \leq (C, D)$ if and only if $A \subseteq C$ and $B \supseteq D$, and the involution is simply $(A, B)^* = (B, A)$. For $(A, B) \in \vec{U}$, we call $A \cap B$ the *separator* of (A, B). It is an *a-b-separator* if $a \in A \setminus B$ and $b \in B \setminus A$. We call $A \cap B$ a *minimal separator* if there exist $a \in A \setminus B$ and $b \in B \setminus A$ for which $A \cap B$ is an inclusion-minimal a-b-separator.

Recall that a *hole* in a graph is an induced cycle on more than three vertices. A graph is *chordal* if it has no holes.

Theorem 18 (Dirac [15])**.** *A graph is chordal if and only if every minimal separator is a clique.*

Let $\vec{S}$ be the set of all $(A, B) \in \vec{U}$ for which $G[A \cap B]$ is a clique. We call these the *clique separations*. Note that $\vec{S}$ is closed under involution and therefore a separation system. To avoid trivialities, we will assume that the graph G is not itself a clique. In particular, this implies that $\vec{S}$ contains no degenerate elements.

Lemma 19. *Let $s, t \in S$. At least three of the four corners of s and t are again in $\vec{S}$. In particular, $\vec{S}$ is submodular.*

Proof. Let $\vec{s} = (A, B)$ and $\vec{t} = (C, D)$. Since $G[A \cap B]$ is a clique and (C, D) is a separation, we must have $A \cap B \subseteq C$ or $A \cap B \subseteq D$, without loss of generality $A \cap B \subseteq C$. Similarly, it follows that $C \cap D \subseteq A$ or $C \cap D \subseteq B$; we assume the former holds. For each corner other than $\vec{s} \wedge \vec{t} = (A \cap C, B \cup D)$, the separator is a subset of either $A \cap B$ or $C \cap D$ and therefore the subgraph it induces is a clique. This proves our claim. □

Suppose that the graph G contains a hole H. Then for every $(A, B) \in \vec{S}$, either $H \subseteq A$ or $H \subseteq B$. In this way, every hole H induces an orientation

$$O_H := \{(A, B) \in \vec{S} \colon H \subseteq B\}$$

of $\vec{S}$. We now describe these orientations as tangles over a suitable set of stars.

Let $\mathcal{F} \subseteq 2^{\vec{U}}$ be the set of all sets $\{(A_1, B_1), \dots (A_n, B_n)\} \subseteq \vec{U}$ for which $G[\bigcap B_i]$ is a clique (note that the graph without any vertices is a clique). As usual, we denote by $\mathcal{F}^*$ the set of all elements of $\mathcal{F}$ which are stars.

Theorem 20. *Let O be an orientation of S. Then the following are equivalent:*

(i) *O is an $\mathcal{F}^*$-tangle.*

(ii) *O is an $\mathcal{F}$-tangle.*

(iii) *There exists a hole H with $O = O_H$.*

It is easy to see that every orientation O_H induced by a hole H is an $\mathcal{F}$-tangle. To prove that, conversely, every $\mathcal{F}$-tangle is induced by a hole, we use Theorem 18 and an easy observation about clique-separators, Lemma 21 below. The proof that every $\mathcal{F}^*$-tangle is already an $\mathcal{F}$-tangle, the main content of Lemma 22 below, is similar to the proof of Lemma 17, but some care is needed to keep track of the separators of two crossing separations.

For a set $\tau \subseteq \vec{U}$, let $J(\tau) := \bigcap_{(A,B)\in\tau} B$ be the intersection of all the right sides of separations in τ, where $J(\emptyset) := V(G)$.

Lemma 21. *Let τ be a set of clique separations, $J = J(\tau)$ and $K \subseteq J$. Let $a, b \in J \setminus K$. If K separates a and b in $G[J]$, then it separates them in G.*

Proof. We prove this by induction on $|\tau|$, the case $\tau = \emptyset$ being trivial. Suppose now $|\tau| \geq 1$ and let $(X, Y) \in \tau$ arbitrary. Put $\tau' := \tau \setminus \{(X, Y)\}$ and $J' := J(\tau')$. Note that $J = J' \cap Y$. Let $G' := G[J']$ and $(X', Y') := (X \cap J', Y \cap J')$.

Then $K \subseteq J'$ and $a, b \in J' \setminus K$. Suppose K did not separate a and b in G' and let $P \subseteq J'$ be an induced a-b-path avoiding K. Since $G'[X' \cap Y']$ is a clique, P has at most two vertices in $X' \cap Y'$ and they are consecutive vertices along P. As $a, b \in Y'$ and (X', Y') is a separation of G', it follows that $P \subseteq Y'$. But then K does not separate a and b in $J = J' \cap Y$, contrary to our assumption.

Hence K separates a and b in G'. By inductive hypothesis applied to τ', it follows that K separates a and b in G. $\square$

Lemma 22. *Every $\mathcal{F}^*$-tangle is an $\mathcal{F}$-tangle and a regular profile.*

Proof. Let P be an $\mathcal{F}^*$-tangle. It is clear that P contains no co-small separation, since $\{(V, A)\} \in \mathcal{F}^*$ for every co-small $(V, A) \in \vec{S}$. Since P is consistent, it follows that P is in fact down-closed.

We now show that P is a profile. Let $(A, B), (C, D) \in P$ and assume for a contradiction that $(E, F) := ((A, B) \vee (C, D))^* \in P$. Recall that either $C \cap D \subseteq A$ or $C \cap D \subseteq B$.

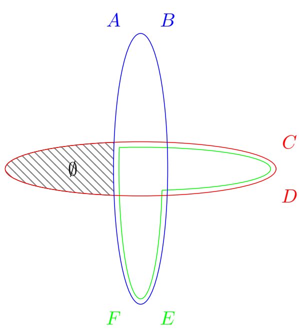

Figure 2: The case $C \cap D \subseteq B$

Suppose first that $C \cap D \subseteq B$; this case is depicted in Figure 2. Let $(X, Y) := (A, B) \wedge (D, C)$ and note that $X \cap Y \subseteq A \cap B$, so that $(X, Y) \in \vec{S}$. It follows from the consistency of P that $(X, Y) \in P$. Let $\tau := \{(C, D), (E, F), (X, Y)\}$ and observe that $\tau \subseteq P$ is a star. However

$$J(\tau) = D \cap (A \cup C) \cap (B \cup C) = (D \cap B) \cap (A \cup C),$$

which is the separator of (E, F). Since $(E, F) \in \vec{S}$, $G[J(\tau)]$ is a clique, thereby contradicting the fact that P is an $\mathcal{F}^*$-tangle.

Suppose now that $C \cap D \subseteq A$. Let $(X, Y) := (B, A) \wedge (C, D)$ and note that $X \cap Y \subseteq A \cap B$, so that $(X, Y) \in \vec{S}$. Since P is down-closed, it follows that $(X, Y) \in P$. Therefore $\tau := \{(A, B), (E, F), (X, Y)\} \subseteq P$. But τ is a star and

$$J(\tau) = B \cap (A \cup C) \cap (A \cup D) = B \cap (A \cup (C \cap D)) = B \cap A,$$

and so $G[J(\tau)]$ is a clique, which again contradicts our assumption that P is an $\mathcal{F}^*$-tangle. This contradiction shows that P is indeed a profile.

We now prove that for any $\tau \subseteq P$ there exists a star $\sigma \subseteq P$ with $J(\sigma) = J(\tau)$. It follows then, in particular, that P is an $\mathcal{F}$-tangle.

Given $\tau \subseteq P$, choose $\sigma \subseteq P$ with $J(\sigma) = J(\tau)$ so that $d(\sigma)$, the number of crossing pairs of elements of σ, is minimum and, subject to this, σ is inclusion-minimal. Then σ is an antichain: If $(A, B) \leq (C, D)$ and both $(A, B), (C, D) \in \sigma$, then $\sigma' := \sigma \setminus \{(A, B)\}$ satisfies $J(\sigma') = J(\sigma)$, thus violating the minimality of σ. Since $\sigma \subseteq P$ and P is consistent, no two elements of σ point away from each other. Therefore, any two nested elements of σ point towards each other. To verify that σ is a star, it suffices to check that σ is nested.

Assume for a contradiction that σ contained two crossing separations (A, B) and (C, D). If $(E, F) := (A, B) \vee (C, D) \in \vec{S}$, obtain σ' from σ by deleting (A, B) and (C, D) and adding (E, F). We have seen above that P is a profile, so $\sigma' \subseteq P$. By Lemma 8, every element of $\sigma \setminus \{(A, B), (C, D)\}$ that is nested with both (A, B) and (C, D) is also nested with (E, F). Since σ' misses the crossing pair $\{(A, B), (C, D)\}$, it follows that $d(\sigma') < d(\sigma)$. But $J(\sigma') = J(\sigma)$, contradicting the minimality of σ.

Hence it must be that $(E, F) \notin \vec{S}$, so $A \cap B \not\subseteq C$ and $C \cap D \not\subseteq A$. Therefore $(X, Y) := (A, B) \wedge (D, C) \in \vec{S}$. Let $\sigma' := (\sigma \setminus \{(A, B)\}) \cup \{(X, Y)\}$. Note that $(X, Y) \leq (A, B) \in P$, so $\sigma' \subseteq P$. Moreover $Y \cap D = (B \cup C) \cap D = B \cap D$, since $C \cap D \subseteq B$. Therefore $J(\sigma') = J(\sigma)$. As mentioned above, any $(U, W) \in \sigma \setminus \{(A, B)\}$ that is nested with (A, B) satisfies $(A, B) \leq (W, U)$. Therefore $(X, Y) \leq (A, B) \leq (W, U)$, so (X, Y) is also nested with (U, W). It follows that $d(\sigma') < d(\sigma)$, which is a contradiction. This completes the proof that σ is nested and therefore a star. $\square$

Proof of Theorem 20. (i) $\to$ (ii): See Lemma 22.

(ii) $\to$ (iii): Let O be an $\mathcal{F}$-tangle and $J := J(O)$. We claim that there is a hole H of G with $H \subseteq J$. Such a hole then trivially satisfies $O_H = O$.

Assume there was no such hole, so that $G[J]$ is a chordal graph. Since O is $\mathcal{F}$-avoiding, $G[J]$ itself cannot be a clique, so there exists a minimal set $K \subseteq J$ separating two vertices $a, b \in J \setminus K$ in $G[J]$. By Theorem 18, K induces a clique in G. By Lemma 21, K separates a and b in G, so there exists a separation $(A, B) \in \vec{S}$ with $A \cap B = K$, $a \in A \setminus B$ and $b \in B \setminus A$. As O orients $\vec{S}$, it must contain one of $(A, B), (B, A)$, say without loss of generality $(A, B) \in O$. But then $J \subseteq B$, contrary to $a \in J$. This proves our claim.

(iii) $\to$ (i): We have $H \subseteq J(O_H)$, so $J(O)$ does not induce a clique. Since every star $\sigma \subseteq O$ has $J(O) \subseteq J(\sigma)$ there is no star $\sigma \subseteq O$ such that $G[J(\sigma)]$ is a clique, and so O is $\mathcal{F}^*$-avoiding. Furthermore, O_H is clearly consistent, and so O is an $\mathcal{F}^*$-tangle. $\square$

The upshot of Theorem 20 is that a hole in a graph, although a very concrete substructure, can be regarded as a tangle. This is in line with our general narrative, set forth e.g. in [12, 6, 7], that tangles arise naturally in very different contexts, and underlines the expressive strength of abstract separation systems and tangles.

What does our abstract theory then tell us about the holes in a graph? The results we will derive are well-known and not particularly deep, but it is nonetheless remarkable that the theory of abstract separation systems, emanating from the theory of highly connected substructures of a graph or matroid, is able to express such natural facts about holes.

Firstly, by Lemma 22, every hole induces a profile of S. Hence Theorem 6 applies and yields a nested set $\mathcal{N}$ of clique-separations distinguishing all holes which can be separated by a clique. This is similar to, but not the same as, the decomposition by clique separators of Tarjan [22]: the algorithm in [22] essentially produces a maximal nested set of clique separations and leaves 'atoms' that do not have any clique separations, whereas our tree set merely distinguishes the holes and leaves larger pieces that might allow further decomposition.

Secondly, we can apply Theorem 4 to find the structure dual to the existence of holes. It is clear that $\mathcal{F}^*$ is standard, since $\mathcal{F}^*$ contains $\{(V, A)\}$ for every $(V, A) \in \vec{S}$.

Lemma 23. *$\mathcal{F}^*$ is closed under shifting.*

Proof. Let $(X, Y) \in \vec{S}$ emulate a non-trivial $(U, W) \in \vec{S}$ with $\{(W, U)\} \notin \mathcal{F}^*$, let $\sigma = \{(A_i, B_i) \colon 0 \le i \le n\} \subseteq \vec{S}$ with $\sigma \in \mathcal{F}^*$ and $(U, W) \le (A_0, B_0)$. Then

$$\sigma' := \sigma^{(X,Y)}_{(A_0,B_0)} = \{(A_0 \cup X, B_0 \cap Y)\} \cup \{(A_i \cap Y, B_i \cup X) \colon 1 \le i \le n\}.$$

By Lemma 10, $\sigma' \subseteq \vec{S}$ is a star. We need to show that $G[J(\sigma')]$ is a clique.

Let $(A, B) := \bigvee_{i \ge 1} (A_i, B_i)$ and note that $(A, B) \le (B_0, A_0)$, since σ is a star. Then

$$(B, A) \wedge (V, B_0) = (B, B_0) \in \vec{U}.$$

But $G[B \cap B_0] = G[J(\sigma)]$ is a clique, so in fact $(B, B_0) \in \vec{S}$. Since $(U, W) \le (A_0, B_0) \le (B, A)$, we see that $(U, W) \le (B, B_0)$. As (X, Y) emulates (U, W) in $\vec{S}$, we find that $(E, F) := (X, Y) \vee (B, B_0) \in \vec{S}$. It thus follows that

$$J(\sigma') = (X \cup B) \cap (Y \cap B_0) = E \cap F$$

and so $G[J(\sigma')]$ is indeed a clique. Therefore $\sigma' \in \mathcal{F}^*$. $\square$

Theorem 24. *Let G be a graph. Then the following are equivalent:*

(i) *G has a tree-decomposition in which every part is a clique.*

(ii) *There exists an S-tree over $\mathcal{F}^*$.*

(iii) *S has no $\mathcal{F}^*$-tangle.*

(iv) *G is chordal.*

Proof. (i) $\to$ (ii): Let $(T, \mathcal{V})$ be a tree-decomposition of G in which every part is a clique. For adjacent $s, t \in T$, let $T_{s,t}$ be the component of $T - st$ containing t and let $V_{s,t}$ be the union of all V_u with $u \in T_{s,t}$. Define $\alpha : \vec{E}(T) \to \vec{U}$ as $\alpha(s,t) := (V_{t,s}, V_{s,t})$. Then $\alpha(s,t) = \alpha(t,s)^*$. The separator of $\alpha(s,t)$ is $V_s \cap V_t$, which is a clique by assumption. Hence (T, α) is in fact an S-tree. It is easy to see that $\alpha(F_t)$ is a star for every $t \in T$ and that $J(\alpha(F_t)) = V_t$. Therefore (T, α) is an S-tree over $\mathcal{F}^*$.

(ii) $\to$ (i): Given an S-tree (T, α) over $\mathcal{F}^*$, define $V_t := J(\alpha(F_t))$ for $t \in T$. It is easily verified that $(T, \mathcal{V})$ is a tree-decomposition of G. Each V_t is then a clique, since $\alpha(F_t) \in \mathcal{F}$.

(ii) $\leftrightarrow$ (iii): Follows from Theorem 4, since $\mathcal{F}^*$ is standard for $\vec{S}$ and closed under shifting by Lemma 23.

(iii) $\leftrightarrow$ (iv): Follows from Theorem 20. $\square$

The equivalence of (i) and (iv) is a well-known characterization of chordal graphs that goes back to a theorem Gavril [19] which identifies chordal graphs as the intersection graphs of subtrees of a tree.

5.3 Tangle-tree duality in cluster analysis

Let us now apply Theorem 4 to a generic scenario in cluster analysis [6, 7], where V is thought of as a data set, S is a set of certain 'natural' bipartitions of V, and we are interested in certain $\mathcal{F}$-tangles as 'clusters'. The idea is that clusters should be described by these $\mathcal{F}$-tangles in the same way as the vertex set of a large grid in a graph is captured by the graph tangle τ it induces: although every oriented separation $\vec{s}$ in τ points to most of the vertices of the grid, the cluster can be 'fuzzy' in that these are not the same points for every $\vec{s} \in \tau$. Indeed, there need not be a single vertex to which all the $\vec{s} \in \tau$ point.

To mimic this idea, we want to choose $\mathcal{F}$ so that, whenever we consider just a few separations in S, any $\mathcal{F}$-tangle τ of S must orient these so that they all point to at least some m (say) points in V, while we do not require that the intersection of all the sets B for $(A, B) \in \tau$ must be large (or even non-empty).

Formally, then, let $\vec{U}$ be the universe of all oriented bipartitions (A, B) of some non-empty set V, including $(\emptyset, V)$ and $(V, \emptyset)$, with $\overleftarrow{s} = (B, A)$ for $\vec{s} = (A, B)$ and $(A, B) \wedge (C, D) := (A \cap C, B \cup D)$, and let $\vec{S} \subseteq \vec{U}$ be any submodular separation system in $\vec{U}$. Let $1 \leq m \in \mathbb{N}$ and $2 \leq n \in \mathbb{N} \cup \{\infty\}$. For these m and n, define

$$\mathcal{F}_m := \{\, F \subseteq \vec{U} : \big| \bigcap_{(A,B) \in F} B \,\big| < m \,\}$$

and

$$\mathcal{F}_m^n := \{\, F \in \mathcal{F}_m : |F| < n \,\}.$$

There is only one small separation in $\vec{U}$, the separation $(\emptyset, V)$. Hence regardless of what S may be, it has no trivial separation other than $(\emptyset, V)$. Since $\{(V, \emptyset)\} \in \mathcal{F}_m^n$ for all m and n, this makes $\mathcal{F}_m^n$ standard for S.

Recall that, for any $\mathcal{F} \subseteq 2^{\vec{S}}$, we write $\mathcal{F}^*$ for the set of stars in $\mathcal{F}$.

Lemma 25. *Let $1 \le m \in \mathbb{N}$ and $2 \le n \in \mathbb{N} \cup \{\infty\}$. For $\mathcal{F} = \mathcal{F}_m^n$, the $\mathcal{F}^*$-tangles of S are precisely its $\mathcal{F}$-tangles.*

Proof. Since $\mathcal{F}^* \subseteq \mathcal{F}$, it is clear that $\mathcal{F}$-tangles are $\mathcal{F}^*$-tangles. To show the converse, suppose there is an $\mathcal{F}^*$-tangle τ that fails to be an $\mathcal{F}$-tangle, because it contains some $F \in \mathcal{F}$ as a subset.

Clearly $F \in \mathcal{F} \setminus \mathcal{F}^*$, so F contains two crossing separations $\vec{r}$ and $\vec{s}$. Since $\vec{S}$ is submodular, one of their opposite corners $\vec{r} \wedge \overleftarrow{s}$ and $\vec{s} \wedge \overleftarrow{r}$ lies in $\vec{S}$; let us assume $\vec{r'} := \vec{r} \wedge \overleftarrow{s}$ does. Since $\vec{r'} \le \vec{r} \in \tau$, the consistency of τ implies that $\vec{r'}$ lies in τ (rather than its inverse $\overleftarrow{r'}$). Indeed, this follows from the definition of consistency if $r' \neq r$. But if $r' = r$ then by $\vec{r'} \le \vec{r}$ either $\vec{r'} = \vec{r} \in \tau$ as desired, or $\vec{r'} < \vec{r} = \overleftarrow{r'}$ with $\vec{r'}$ small but $\overleftarrow{r'} \in \tau$. Since $(\emptyset, V)$ is the only small separation in $\vec{S}$ and is in fact trivial, the consistency of τ once more implies that $\vec{r'} \in \tau$.

Let F' be obtained from F by replacing $\vec{r}$ with $\vec{r'}$. Note that $\bigcap_{(A,B)\in F} B = \bigcap_{(A',B')\in F'} B'$ has remained unchanged: although we replaced the set B from $\vec{r} = (A, B)$ in the first intersection with the bigger set B' from $\vec{r'} = (A', B')$ in the second, the additional $B' \setminus B$ has empty intersection with the set D from $\vec{s} = (C, D)$, and therefore does not increase the second intersection. Hence our assumption of $F \in \mathcal{F} = \mathcal{F}_m^n$ implies that also $F' \in \mathcal{F}_m^n = \mathcal{F}$.

Note that while $\vec{r}$ and $\vec{s}$ crossed, $\vec{r'}$ and $\vec{s}$ are nested; indeed, $\{\vec{r'}, \vec{s}\}$ is a star. Moreover, replacing an element of this star by a smaller separation will yield another star; in particular, it cannot result in another pair of crossing separations. This means that iterating the above uncrossing procedure of replacing in F an element $\vec{r} \in \tau$ with a smaller separation $\vec{r'} \in \tau$ in a way that keeps F in $\mathcal{F}$ will end after at most $\binom{|F|}{2}$ steps: for every 2-set $\{\vec{r}, \vec{s}\} \subseteq F$ we will consider only once in this iterated process a pair $\{\vec{r'}, \vec{s'}\}$ where $\vec{r'}$ is either $\vec{r}$ or a replacement of $\vec{r}$, and $\vec{s'}$ is either $\vec{s}$ or a replacement of $\vec{s}$.

Since the above process turns every pair of crossing separations from F into a 2-star of separations, and a set of separations is a star as soon as all its 2-subsets are stars, the set we turn F into will be a star in $\mathcal{F}$, an element of $\mathcal{F}^*$. As it will also still be a subset of τ, this contradicts our assumption that τ is an $\mathcal{F}^*$-tangle. $\square$

Lemma 26. *Let $1 \le m \in \mathbb{N}$ and $2 \le n \in \mathbb{N} \cup \{\infty\}$. The set $\mathcal{F}^*$ of stars in $\mathcal{F} = \mathcal{F}_m^n$ is closed under shifting.*

Proof. Suppose that $\vec{s} \in \vec{S}$ emulates in $\vec{S}$ some nontrivial $\vec{r}$ not forced by $\mathcal{F}$. We have to show that for every star $\sigma \subseteq \vec{S} \setminus \{\overleftarrow{r}\}$ with $\sigma \in \mathcal{F}^*$ and every $\vec{x} \in \sigma$ with $\vec{x} \ge \vec{r}$ we have $\sigma' := \sigma^{\vec{s}}_{\vec{x}} \in \mathcal{F}^*$.

Let $\vec{s} = (U, W)$, and for $(A, B) \in \sigma$ write $(A', B') \in \sigma'$ for the separation that (A, B) shifts to: if $(A, B) = \vec{x}$ then $(A', B') := (A \cup U, B \cap W)$, while if $(A, B) \in \sigma \setminus \{\vec{x}\}$ then $(A', B') := (A \cap W, B \cup U)$. From these explicit representations of the elements of σ' it is clear that

$$\bigcap_{(A',B') \in \sigma'} B' = \bigcap_{(A,B) \in \sigma} B' \subseteq \bigcap_{(A,B) \in \sigma} B,$$

since $B' \setminus B \subseteq U$ for every $(A, B) \in \sigma \setminus \{\vec{x}\}$ while $U \cap B' = \emptyset$ for $(A, B) = \vec{x}$, so that the overall intersection of all the B' equals that of all the B. And since these sets did not change, nor did their cardinality: as $\sigma \in \mathcal{F} = \mathcal{F}^n_m$ we also have $\sigma' \in \mathcal{F}^n_m = \mathcal{F}$. By Lemma 10, this implies $\sigma' \in \mathcal{F}^*$ as desired. $\square$

Together with Lemmas 25 and 26, Theorem 4 implies our last main theorem:

Theorem 5. *Let $\vec{S} \subseteq \vec{U}$ be a submodular separation system, let $1 \le m \in \mathbb{N}$ and $2 \le n \in \mathbb{N} \cup \{\infty\}$, and let $\mathcal{F} = \mathcal{F}^n_m$. Then exactly one of the following two statements holds:*

(i) *S has an $\mathcal{F}$-tangle;*

(ii) *There exists an S-tree over $\mathcal{F}^*$.* $\square$

5.4 Phylogenetic trees from tangles of circle separations

Finally, let us describe an application of Theorem 1 and Theorem 5 in biology. Let V be a set, which we think of as a set of species, or of (possibly unknown) organisms, or of DNA samples. Our aim is to find their Darwinian 'tree of life': a way of dividing V recursively into ever-smaller subsets so that the leaves of this division tree correspond to the individual elements of V.

This tree can be formalized by starting with a root node labelled $\emptyset$ at level 0, a unique node labelled V at level 1, and then recursively adding children to each node, labelled $A \subseteq V$, say, corresponding to the subsets of A into which A is divided and labelling these children with the subsets to which they correspond. Furthermore, let us label the edges from level $k - 1$ to k by k.

Every edge e of this tree naturally defines a bipartition of V, to which we assign the label k of e as it 'order'. Species that are fundamentally different are then separated by a bipartition of low order, while closely related but distinct species are only separated by bipartitions of higher order.

If we draw that tree in the plane, the leaves – and hence V – will be arranged in a circle, C say. The bipartitions $\{A, B\}$ of V defined by the edges of the tree are given by *circle separations*: A and B are covered by disjoint half-open segments of C whose union is C.

The way in which this tree is found in practice is roughly as follows. One first defines a metric on V in which species u, v are far apart if they differ in many respects. Then one applies some clustering algorithm, such as starting with the singletons $\{v\}$ for $v \in V$ as tiny clusters and then successively amalgamating close clusters (in terms of this distance function) into bigger ones. The bipartitions of V corresponding to pitching a single cluster against the rest of V will be nested and can therefore be represented as edge-separations of a tree as above, which is then output by the algorithm. If we draw the tree in the plane, the bipartitions then become circle separations of V as earlier.

A known problem with this approach is that, since every cluster found in this process defines a bipartition of V that ends up corresponding to an edge of the final tree, inaccuracies in the clustering process immediately affect this tree in an irreversible way. Bryant and Moulton [1] have suggested a more careful clustering process which produces not necessarily a tree but an outerplanar graph G on V, together with a set of particularly important bipartitions of V that are not necessarily nested but are still circle separations of V with respect to the outer face boundary of G. The task then is to select from the set S of these bipartition a nested subset that defines the desired phylogenetic tree, or perhaps to generate such a set from S in some other suitable way such as adding corners of crossing separations already selected.

This is where tangles can help: in a general way, but also in a rather specific way that finds the desired nested set from a set of circular separations of V.

Let us just briefly indicate the general way in which tangles can be used for finding phylogenetic trees in a novel way [6], without the need for any distance-based clustering. As input we need a collection of subsets of V to be used as similarity criteria, such as the set of species $v \in V$ that can fly or lay eggs, the set of DNA molecules that have base T in position 137, or the set of those organisms that respond to some test in a certain way. Then we define an order function on all the bipartitions of V, assigning low order to those bipartitions that do not cut accross many of our criteria sets so as to split them nearly in half. For example, we might count for $s = \{A, B\}$ the number of triples (a, b, c) such that $a \in A$ and $b \in B$ and both a and b satisfy (are elements of) the criterion c.

This order function is easily seen to be submodular on the universe $\vec{U}$ of all oriented bipartitions of V [7]. For suitable $\mathcal{F}$ whose $\mathcal{F}$-tangles are profiles, e.g. the $\mathcal{F} = \mathcal{F}_m^n$ with $n > 3$ considered in Section 5.3, we can then compute the (canonical) tree of tangles as in [11], or an S-tree over $\mathcal{F}$ as in [12, 13] if there are no tangles. In the first case, the tree of tangles for $\mathcal{F}_2^n$ has a good claim to be the phylogenetic tree for the species in V, see [6].

In the concrete scenario of [1], we further have the following specific application of tangles as studied in this paper. Let S be the set of circle separations of V, taken with respect to its circular ordering found by the current algorithm of [1]. We can now define $\mathcal{F}$-tangles on this S just as we did earlier on the set of all bipartitions of V, and consider the same order function as earlier.

This order function is not, however, submodular on our restricted set S: this would require that S is not just a separation system but a universe of separations, i.e., that corner separations of elements of S are again in S – which is not always the case. In particular, we do not get a tree-of-tangles theorem for S from [11], or a tangle-tree duality theorem from [12, 13].[3]

[3]This is not to say that no submodular order function on S exists that returns the sets S_k

However, we do get a tree-of-tangles theorem for S as a corollary of Theorem 1, and a tangle-tree duality theorem as a corollary of Theorem 5. For this we need the following easy lemma. Let $\vec{U}$ be the universe of all oriented bipartitions of V (see Section 5.3), equipped with any submodular order function under which $\{\emptyset, V\}$ has order 0. Let S be the set of all circle separations of V with respect to some fixed cyclic ordering of V.

Lemma 27. *For every $k > 0$, the set S_k of all circle separations of order $< k$ is submodular.*

Proof. Consider two oriented circle separations $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$ of V. Clearly, $\vec{r} \vee \vec{s}$ is again a circle separation unless the circle segments representing A and C are disjoint and both segments that join them on the circle meet V. In that case, however, the union of the segments representing B and D is the entire circle, so $\vec{r} \wedge \vec{s} = (\emptyset, V)$, which is a circle separation.

As $(\emptyset, V)$ has order $0 < k$ by assumption, and our order function is submodular on the set of all bipartitions of V, this implies that S_k is submodular: given $\vec{r}, \vec{s} \in \vec{S_k}$, either both $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{r}$ are in $\vec{S}$ and hence one of them is in $\vec{S_k}$, or one of them is $(\emptyset, V)$ or $(V, \emptyset)$ and therefore in $\vec{S_k}$. □

Consider any $1 \leq m \in \mathbb{N}$ and $3 < n \in \mathbb{N} \cup \{\infty\}$. Let $\mathcal{F} = \mathcal{F}_m^n$ be as defined in Section 5.3. Here is our first tree-of-tangles theorem for circle separations:

Theorem 28. *For every $k > 0$, the set S_k of circle separations of V of order $< k$ contains a tree set of separations that distinguishes all the $\mathcal{F}$-tangles of S_k.*

Proof. In order to apply Theorem 1, we have to show that all $\mathcal{F}$-tangles of S_k are abstract tangles, i.e., that they contain no triple $(\vec{r}, \vec{s}, \vec{t})$ with $\vec{r} \vee \vec{s} \vee \vec{t} = (V, \emptyset)$ (which is the unique co-small separation in $\vec{U}$). But any such triple lies in $\mathcal{F}$ for the values of m and n we specified, so no $\mathcal{F}$-tangle of S_k contains it.

Lemma 27 and Theorem 1 thus imply the result. □

Applying very recent work of Elbracht, Kneip and Teegen [16], we can unify the assertions of Theorem 28 over all k, as follows. Let us say that an element s of S *distinguishes* two orientations ρ, τ of subsets of S if these orient s differently, i.e., if s has orientations $\vec{s} \in \rho$ and $\overleftarrow{s} \in \tau$. If s has minimum order amongst the separations in S that distinguish ρ from τ, we say that s distinguishes ρ and τ *efficiently*. Similarly, a set $T \subseteq S$ *distinguishes* a set $\mathcal{T}$ of orientations of subsets of S *efficiently* if for every pair of distinct $\rho, \tau \in \mathcal{T}$ there exists an $s \in T$ that distinguishes ρ from τ efficiently.

Orientations ρ and τ as above are called *distinguishable* if they are distinguished by some $s \in S$. Note that orientations ρ of S_k and τ of S_ℓ for $k \leq \ell$ are indistinguishable if and only if $\rho = \tau \cap \vec{S_k}$.

Here is our tree-of-tangles theorem for circle separations of mixed order:

Theorem 29. [16] *The set S of all circle separations of V contains a tree set that efficiently distinguishes all the distinguishable $\mathcal{F}$-tangles of subsets S_k of S.*

---

we are interested in as sets $S_{k'}$ for some other k'. One can indeed construct such a function, but it is neither obvious nor natural.

Elbracht, Kneip and Teegen [16] showed that the tree set in Theorem 29 can in fact be chosen *canonical*, i.e., so that every separation system automorphism (see [8]) of $\vec{S}$ acts on $\vec{T}$ as a set of automorphisms of T.

Finally, our tangle-tree duality theorem for circle separations:

Theorem 30. *For every $k > 0$, the set S_k of circle separations of V of order $< k$ satisfies exactly one of the following two assertions:*

(i) *S_k has an $\mathcal{F}$-tangle;*

(ii) *There exists an S_k-tree over $\mathcal{F}^*$.*

Proof. Apply Lemma 27 and Theorem 5. □

Acknowledgement

Geoff Whittle has informed us that his student Jasmine Hall has obtained explicit proofs of Theorems 28 and 30 by re-working the theory of graph and matroid tangles given in [20, 21].

References

[1] David Bryant and Vincent Moulton. Neighbor-net: an agglomerative method for the construction of phylogenetic networks. *Mol. Biol. Evol.*, 21(2):255–265, 2004.

[2] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs I. Existence and algorithms. *J. Combin. Theory Ser. B*, 116:1–24, 2016.

[3] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs II. Essential parts. *J. Combin. Theory Ser. B*, 118:268–283, 2016.

[4] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. Connectivity and tree structure in finite graphs. *Combinatorica*, 34(1):1–35, 2014.

[5] J. Carmesin and P. Gollin. Canonical tree-decompositions of a graph that display its k-blocks. *J. Combin. Theory Ser. B*, 122:1–20, 2017.

[6] R. Diestel. Tangles in the empirical sciences: indirect clustering by separation. In preparation.

[7] R. Diestel. Tangles in the social sciences: a new mathematical model to identify types and predict behaviour. In preparation.

[8] R. Diestel. Abstract separation systems. *Order*, 35:157–170, 2018.

[9] R. Diestel. Tree sets. *Order*, 35:171–192, 2018.

[10] R. Diestel, Ph. Eberenz, and J. Erde. Duality theorem for blocks and tangles in graphs. *SIAM J. Discrete Math.*, 31(3):1514–1528, 2017.

[11] R. Diestel, F. Hundertmark, and S. Lemanczyk. Profiles of separations: in graphs, matroids, and beyond. arXiv:1110.6207, to appear in *Combinatorica.*

[12] R. Diestel and S. Oum. Tangle-tree duality in graphs, matroids and beyond. arXiv:1701.02651. To appear in *Combinatorica.*

[13] R. Diestel and S. Oum. Tangle-tree duality in abstract separation systems. Submitted, arXiv:1701.02509, 2017.

[14] R. Diestel and G. Whittle. Tangles and the Mona Lisa. arXiv:1603.06652.

[15] G. A. Dirac. On rigid circuit graphs. *Abh. Math. Sem. Univ. Hamburg*, 25:71–76, 1961.

[16] C. Elbracht, J. Kneip, and M. Teegen. Trees of tangles in abstract separation systems. In preparation.

[17] Joshua Erde. Refining a tree-decomposition which distinguishes tangles. *SIAM J. Discrete Math.*, 31:1529–1551, 2017.

[18] Joshua Erde. A unified treatment of linked and lean tree-decompositions. *J. Combin. Theory Ser. B*, 130:114–143, 2018.

[19] F. Gavril. The intersection graphs of subtrees in trees are exactly the chordal graphs. *J. Combinatorial Theory Ser. B*, 16:47–56, 1974.

[20] J. Geelen, B. Gerards, N. Robertson, and G. Whittle. Obstructions to branch-decomposition of matroids. *J. Combin. Theory (Series B)*, 96:560–570, 2006.

[21] N. Robertson and P.D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)*, 52:153–190, 1991.

[22] R.E. Tarjan. Decomposition by clique separators. *Discrete mathematics*, 55(2):221–232, 1985.