Structural submodularity and tangles in abstract separation systems

Reinhard Diestel, Joshua Erde and Daniel Weißauer

Abstract

We prove a tangle-tree theorem and a tangle 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 the two theorems 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 [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 13, 14] 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 [11] or the social sciences [9]. 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 [17]. One of these is the *tangle-tree* 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 [17], the *tangle 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 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.

A notion of separation does come naturally with many structures. 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 [5].

The second step, then, was to generalize the proofs of the tangle-tree and the tangle duality theorem to the abstract setting of separation systems. This was done in [8] and [9], respectively.

In order to prove these theorems, or to apply them to concrete cases of abstract separation systems, e.g. as in [10, 11], 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]

With this third step, then, the programme sketched above will be complete:

[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 as separation systems. For separation universes, therefore, submodularity is used with the narrower meaning of being endowed with a submodular order function [5].

we shall have a notion of tangle for very general abstract separation systems, as well as a tangle-tree and a 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.)

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 Section 5.2, 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 [5]; 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 [8, 9, 10] and in [11]. 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* $(\vec{S}, \leq, ^*)$ 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) separation* 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 the set of all separations of *low order* with respect to some order function $|\cdot|\colon \vec{U} \to \mathbb{R}^+$ that is *symmetric* and *submodular*, that is, $|\vec{s}| = |\overleftarrow{s}|$ and $|\vec{s} \vee \vec{t}| + |\vec{s} \wedge \vec{t}| \leq |\vec{s}| + |\vec{t}|$ 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 *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 $\vec{r} \leq \vec{s} \in \vec{S}$, then $\overleftarrow{r} \notin \vec{S}$. 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 asking that three oriented separations in an 'abstract tangle' must never have a co-small supremum; see [5, 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.

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

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 tangle-tree theorem for abstract tangles:

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 duality theorem, 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 $\vec{S}$ consists of the bipartitions of some finite set X, then every $x \in X$ induces an abstract tangle

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

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

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 θ of S *extends* Q if

$Q \cap \vec{S} \subseteq \theta$. It is easy to see that θ extends Q if and only if θ 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 $\vec{r} \leq \vec{s} \in Q$ implies $\vec{r} \in Q$ for all $\vec{r}, \vec{s} \in \vec{U}$.

Here, then, is our refined duality theorem for abstract tangles.

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 $Q = \emptyset$.

3 The tangle-tree theorem

In this section we will prove Theorem 1. In fact, we are going to prove a slightly more general statement. Let $\mathcal{P} := \{\{\vec{s}, \vec{t}, (\vec{s} \vee \vec{t})^*\} \colon \vec{s}, \vec{t} \in \vec{U}\}$. The $\mathcal{P}$-tangles are known as *profiles*.

Theorem 4. *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 5. *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. □

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 6 ([5])**.** *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 4, 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 7. *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. □

Proof of Theorem 4. 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.

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 $\vec{s} \leq \vec{r_1} \in Q$: 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 6, 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 $\vec{t} \leq \vec{u}$ or $\vec{t} \leq \overleftarrow{u}$, since $\vec{r_2} \leq \vec{t}$. It cannot be that $\overleftarrow{u} \leq \vec{t}$, 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

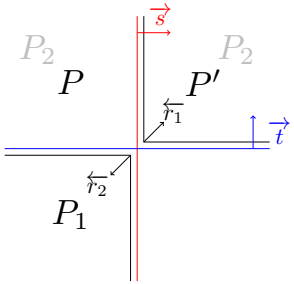

Figure 1: Crossing separations

and s are nested and it follows from Lemma 6 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 $\vec{r_2} \leq \vec{t} \in P_2$: if $\vec{r_2} = \vec{t}$, 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. □

4 Duality

Our goal in this section is to prove Theorem 3. The proof will be an application of a more general duality theorem of Diestel and Oum. We first need to introduce the central notion of separability.

A separation $\vec{s} \in \vec{S}$ *emulates* $\vec{r}$ *in* $\vec{S}$ if $\vec{s} \geq \vec{r}$ and for every $\vec{t} \in \vec{S} \setminus \{\overleftarrow{r}\}$ with $\vec{t} \geq \vec{r}$ 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

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

Lemma 8. *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 $\vec{r} \leq \vec{x} \in \sigma$. 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 $\vec{r} \leq \overleftarrow{y}$. It is clear that for any two distinct $\vec{u}, \vec{v} \in \sigma^{\vec{s}}_{\vec{x}}$ we have $\vec{u} \leq \overleftarrow{v}$, 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 $\vec{r} \leq \vec{t}$ 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 $\vec{r} \leq \vec{t}$. 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 $\vec{r} \leq \vec{t}$. This contradicts our assumption on $\vec{r}$. □

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 $\vec{x} \geq \vec{r}$ 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 $\vec{r_1}, \vec{r_2} \in \vec{S}$ with $\vec{r_1} \leq \vec{r_2}$ and $\{\overleftarrow{r_1}\}, \{\overleftarrow{r_2}\} \notin \mathcal{F}$ 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 9 ([9, 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 $\mathcal{F}$.*

Since the family $\mathcal{F}$ in Theorem 9 is assumed to be a set of stars, we cannot work directly with $\mathcal{T} \cup \mathcal{F}_Q$. We thus keep only those sets in $\mathcal{T}$ which happen to be stars and apply Theorem 9 with $\mathcal{F} = \mathcal{T}_Q$, where $\mathcal{T}_Q := \mathcal{T}^* \cup \mathcal{F}_Q$. As it turns out, this does not make a difference as long as $\vec{S}$ has no degenerate elements, see Lemma 15 below.

We start with a simple observation that will be useful later.

Lemma 10. *Let $\vec{U}$ be a distributive universe of separations. Let $\vec{u}, \vec{v}, \vec{w} \in \vec{U}$. If $\vec{u} \leq \vec{v}$ 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 $\vec{U}$

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

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

$$\overleftarrow{x} \leq \overleftarrow{s} \vee \overleftarrow{t} \leq \vec{s} \wedge \vec{t} \leq \vec{x}. \qquad \square$$

We now prove that $\vec{S}$ is $\mathcal{T}_Q$-separable in a strong sense.

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 strong separability extend verbatim to a separation system within a universe of separations when regarded as a subset of the underlying lattice. The notions of lifting and emulating are of course closely related: If $\vec{s} \in \vec{S}$ lifts $\vec{r} \in \vec{S}$, then $\vec{s}$ emulates $\vec{r}$ in $\vec{S}$. Observe also that $\vec{s}$ pushes $\vec{r}$ if and only if $\overleftarrow{s}$ lifts $\overleftarrow{r}$.

We call a set $\mathcal{F} \subseteq 2^{\vec{U}}$ *closed under shifting* if whenever $\vec{s} \in \vec{S}$ emulates in $\vec{S}$ a non-trivial and non-degenerate $\vec{r} \in \vec{S}$ with $\{\overleftarrow{r}\} \notin \mathcal{F}$, then it does so for $\mathcal{F}$.

The following is immediate from the definitions:

Lemma 11. *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 strongly separable and $\mathcal{F}$ is closed under shifting, then $\vec{S}$ is $\mathcal{F}$-separable.* $\square$

Lemma 12. *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 8 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 10 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 cosmall. 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

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

Virtually all applications of Theorem 9 given in [9] involve a separation system of the form $\vec{S} = \vec{S}_k$ consisting of all separations of order $< k$ within some ambient universe $\vec{U}$ endowed with a symmetric and submodular order function. In most situations, submodularity is only used to ensure that at least one of any two opposite corners of two separations s, t both of order $< k$ again has order $< k$ – which is tantamount to saying that $\vec{S}_k$ is structurally submodular. (Indeed, this fact motivated our abstract notion of submodularity.)

The proof that $\vec{S}_k$ is separable, however – see [10, Lemma 3.4] – requires a more subtle use of the submodularity of the order function: the orders of the two corners are not compared with any fixed value of k, but with the possibly distinct orders of s and t. This kind of argument is naturally difficult, if not impossible, to mimic in our set-up. As a consequence, separability was added as an explicit assumption on the submodular separation system in [7, Theorem 3.9].

However, we can prove that every submodular separation system is in fact separable, thereby showing that this additional assumption may be removed:

Lemma 13. *Let L be a finite lattice and $M \subseteq L$ submodular. Then M is strongly separable.*

(This lemma allows further applications of Theorem 9 beyond the present context of abstract tangles. For instance, we will make use of it in Section 5.2 to show that the separation system of clique separations of a graph is 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 $I(a, b) := \{u \in M \colon a \leq u \leq b\}$ 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. □

Theorem 14. *Let $\vec{S}$ be a submodular separation system in a distributive universe $\vec{U}$ and let $Q \subseteq \vec{U}$ down-closed. Then exactly one of the following holds:*

(i) *There exists a $\mathcal{T}_Q$-tangle of S.*

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

Proof. By Lemmas 12 and 13, $\vec{S}$ is $\mathcal{T}_Q$-separable. Since every trivial element is small and non-degenerate, $\mathcal{T}_Q$ is standard for $\vec{S}$. Hence Theorem 9 applies and yields the desired duality. □

Our original aim was a duality theorem for abstract tangles, not for $\mathcal{T}^*$-tangles. However, as long as $\vec{S}$ contains no degenerate elements, these notions coincide:

Lemma 15. *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 $\vec{v} \vee \vec{w} = \bigvee \sigma$ is co-small, we can apply Lemma 10 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. □

When $\vec{S}$ has no degenerate elements, the abstract tangles extending Q are precisely the $\mathcal{T}^*$-tangles extending Q (by Lemma 15), which are exactly the $\mathcal{T}_Q$-tangles. Therefore, Theorem 3 is an immediate consequence of Theorem 14.

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 [17], but a good deal of the work is done in the setting of connectivity systems. Geelen, Gerards and Whittle [16] 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 $Q := \{\emptyset\} \cup \{\{x\} \colon x \in X\}$ 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 [16], 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 [17, 16].

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 9 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 16 (Dirac [12])**.** *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 17. *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 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 18. *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 16 and an easy observation about clique-separators, Lemma 19 below. The proof that every $\mathcal{F}^*$-tangle is already an $\mathcal{F}$-tangle, the main content of Lemma 20 below, is similar to the proof of Lemma 15, 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 19. *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. □

Lemma 20. *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 $C \cap D \subseteq B$.

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 $(E, F) \in \vec{S}$, the latter is a clique, thereby contradicting the fact that P is an $\mathcal{F}^*$-tangle.

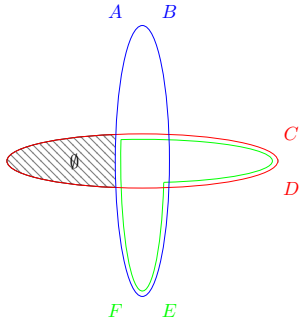

Figure 2: The case $C \cap D \subseteq B$

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

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

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 6, 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 18. (i) $\to$ (ii): See Lemma 20.

(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 16, K induces a clique in G. By Lemma 19, 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 $J(O_H)$ does not induce a clique. $\square$

The upshot of Theorem 18 is that a hole in a graph, although a very concrete substructure, can be regarded as a tangle. This is in line with the idea, set forth in [11], 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.

First, by Lemma 20, every hole induces a profile of S. Hence Theorem 4 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 [18]: the algorithm in [18] 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.

Second, we can apply Theorem 9 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 21. *$\vec{S}$ is $\mathcal{F}^*$-separable.*

Proof. By Lemma 17 and Lemma 13, $\vec{S}$ is strongly separable. We show that $\mathcal{F}^*$ is closed under shifting.

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

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

By Lemma 8, $\sigma' \subseteq \vec{S}$ is a star. We need to show that $J(\sigma')$ is a clique.

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

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

But $B \cap B_0 = J(\sigma)$ is a clique, so in fact $(B, B_0) \in \vec{S}$. Since $(U, W) \leq (A_0, B_0) \leq (B, A)$, we see that $(U, W) \leq (B, B_0)$. 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

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

is indeed a clique. Therefore $\sigma' \in \mathcal{F}^*$. □

Theorem 22. *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) → (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) → (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) ↔ (iii): Follows from Theorem 9, since $\mathcal{F}^*$ is standard for $\vec{S}$ and $\vec{S}$ is $\mathcal{F}^*$-separable by Lemma 21.

(iii) ↔ (iv): Follows from Theorem 18. □

The equivalence of (i) and (iv) is a well-known characterization of chordal graphs that goes back to a theorem Gavril [15] which identifies chordal graphs as the intersection graphs of subtrees of a tree.

References

[1] 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.

[2] 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.

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

[4] 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.

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

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

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

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

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

[10] R. Diestel and S. Oum. Tangle-tree duality in graphs, matroids and beyond. arXiv:1701.02651, 2017.

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

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

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

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

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

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

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

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