

Tangle-tree duality in abstract separation systems

Reinhard Diestel

Mathematisches Seminar, Universität Hamburg

Sang-il Oum

KAIST, Daejeon, 34141 South Korea

January 10, 2017

Abstract

We prove a general width duality theorem for combinatorial structures with well-defined notions of cohesion and separation, such as graphs and matroids. The theorem asserts a duality between the existence of high cohesiveness somewhere local and a global overall tree structure.

We describe cohesive substructures in a unified way in the format of tangles: as orientations of low-order separations satisfying certain consistency axioms. These axioms can be expressed without reference to the underlying structure, such as a graph or matroid, but just in terms of the poset of the separations themselves. This makes it become possible to identify tangles, and apply our tangle-tree duality theorem, in very diverse settings.

Our result implies all the classical duality theorems for width parameters in graph minor theory, such as path-width, tree-width, branch-width or rank-width. It yields new, tangle-type, duality theorems for tree-width and path-width. It implies the existence of width parameters dual to cohesive substructures such as k -blocks, edge-tangles, or given subsets of tangles, for which no width duality theorems were previously known.

Abstract separation systems can be found also in structures quite unlike graphs and matroids. For example, our theorem can be applied to image analysis by capturing the regions of an image as tangles of separations defined as natural partitions of its set of pixels. It could also be applied in pure mathematics, e.g. to separations of compact manifolds.

1 Introduction

There are a number of theorems in the structure theory of sparse graphs that assert a duality between high connectivity present somewhere in the graph and an overall tree structure. For example, a graph either has a large complete minor or a tree-decomposition into torsos of essentially bounded genus, but not both [4, 14]. And it either has a large grid minor or a tree-decomposition into parts of bounded size, but not both [4, 12]. Let us loosely refer to such highly connected substructures of a graph, defined in terms of subsets of its vertices together with some required edges, as *concrete* highly cohesive substructures (*HCSs*).

An example of a concrete HCS for which no dual notion of global tree structure is known is that of a *k-block* [2, 10]: a set of at least k vertices no two of which can be separated in the graph by deleting less than k vertices.

Conversely, there are a number of so-called width parameters for graphs, invariants whose boundedness asserts that the graph has some kind of global tree structure, for which there are no obvious dual concrete HCSs.

Amini, Mazoit, Nisse, and Thomassé [1] addressed this latter problem in a broad way: they showed how to construct, for many width-parameters including all the then known ones, dual concrete HCSs akin to brambles.¹ For each parameter, the existence of such an HCS forces this parameter to be large, and conversely, whenever one of these width parameters is large there exists a concrete bramble-type HCS to witness this.

In one of their seminal papers on graph minors [13], Robertson and Seymour introduced a very different way to capture high cohesion somewhere in a graph, which they call *tangles*. The basic idea behind these is as follows. Given a concrete HCS X in a graph G and a low-order separation,² most of X will lie on one of its two sides: otherwise X could not be highly cohesive.³ In this way X , whatever it is, *orients* each of the low-order separations of G towards one of its sides, the side that contains most of X . These orientations of all the low-order separations will be ‘consistent’ in various ways, since they all point towards X : no two of them, for example, will point away from each other.

Robertson and Seymour [13] noticed that these orientations of all the low-order separations of G captured most of what they needed to know about X . Consequently, they defined a *tangle of order k* in a graph as a way to orient all its separations of order $< k$, consistently in some precise sense not relevant here.

The notion of a tangle brought with it a shift of paradigm in the connectivity theory of graphs [11]: we can now think of a consistent orientation of all the low-order separations of a graph as a ‘highly cohesive substructure’ in its own right: no longer a concrete one, but an *abstract HCS*. Such abstract HCSs, though maybe unfamiliar at first, are often ‘deeper’ than concrete ones, because they pick out only the essential information. But they are also easier to work with: one no longer has to worry about the details, say, of where exactly in the graph a subdivided grid has all its connecting paths. And most importantly, they are able to capture HCSs that are inherently fuzzy. For example, the additional detail that a subdivided grid contains over the tangle it defines is not only superfluous but can be misleading: each individual branch vertex can, and typically will, lie on the *wrong* side of *some* low-order separation, the side that does not contain most of the grid. (Consider, for example, the separation defined by the four neighbours of a given vertex in an actual grid.)

Our first aim in this paper is to do for abstract HCSs in graphs and matroids the converse of what Amini et al. did for concrete ones: to come up with a

¹Brambles are collections of pairwise ‘touching’ connected vertex sets in a graph with no small transversal. They were introduced by Seymour and Thomas [15] as a concrete HCS dual to low tree-width. See [4] for an introduction.

²A *separation* of G is a pair $\{A, B\}$ of vertex sets such that $G = G[A] \cup G[B]$. Its *order* is the number $|A \cap B|$.

³The exact way how to make this precise will depend on the type of HCS that X represents. For example, if X is a complete minor then its branch sets either all meet A or all meet B . If X is a large subdivided grid H , then one of the sets A, B will contain most of its branch vertices: there cannot be many in A as well as many in B , since any two large sets of branch vertices can be joined in H , and hence in G , by more than $|A \cap B|$ disjoint paths if H is large.

unified way of describing abstract HCSs, and then to prove a general duality theorem that would find corresponding tree structures witnessing the possible nonexistence of these HCSs.

Generalizing the specific notion of a tangle from [13], we shall define types of abstract HCSs to be called ‘ \mathcal{F} -tangles’, where \mathcal{F} encodes some particular type of consistency. Thus, an \mathcal{F} -tangle in a graph will be an orientation of all its separations of order $< k$ for some k that are consistent in a sense specified by \mathcal{F} : different notions of consistency will give rise to different sets \mathcal{F} and result in different \mathcal{F} -tangles. But we shall prove one unified duality theorem saying that, for every suitable \mathcal{F} , a given graph either has an \mathcal{F} -tangle or a global tree structure that clearly precludes the existence of an \mathcal{F} -tangle.

Our duality theorem will easily imply the two known tangle-type duality theorems from graph minor theory: the classical Robertson-Seymour one for tangles and branch-width in graphs [13], and its analogue for matroids [8, 13]. This will be shown in detail in [6].

It will also imply new, tangle-type, duality theorems for all the other classical width parameters: for each of these we shall find an \mathcal{F} , encoding some specific type of consistency, such that the graphs where this parameter is large are precisely those with an \mathcal{F} -tangle.⁴ The known duality theorems for these width parameters, in terms of concrete HCSs, will follow from our duality theorem in terms of abstract HCSs, but not conversely. This, too, will be shown in [6].

Our result will also imply duality theorems for k -blocks, the main concrete HCS for which no duality theorem has been known, and for any specified type of classical tangles (rather than all of them). This will be done in [5].

While the study of tangles as abstract HCSs marked a shift of paradigm from the earlier studies of concrete HCSs, there has since been another major shift of paradigm: from concrete to ‘abstract’ separations. Separations in graphs – as well as traditional tangles and their dual branch decompositions – are defined in terms of the graph’s edges. But when we proved our duality theorem for graphs we found that, surprisingly, we needed to know only how these separations relate to each other, not how they relate to the graph which they separate.

Our main result, therefore, is now a duality theorem for *abstract separation systems*. Very roughly, these are partially ordered sets (reflecting the natural partial ordering between separations in graphs and matroids), with an order-reversing involution that reflects the flip $(A, B) \mapsto (B, A)$ of a graph separation.

Both tangles in graphs and their dual tree structure can be expressed in terms of just this partial ordering of their separations. Indeed, the consistency requirement for tangles, that no three ‘small’ sides of its oriented separations shall cover the graph, can be replaced by the requirement that whenever a tangle τ contains two oriented separations, (A, B) and (C, D) say, it also contains their supremum $(A \cup C, B \cap D)$ as long as this is oriented by τ at all, i.e., has order $< k$ if τ is a k -tangle.⁵ And the tree-decompositions or branch-decompositions dual to graph tangles can be described purely in terms of the separations that correspond to the edges of their decomposition trees, where the requirement that these edges form a tree can be replaced by requiring that those separations must be nested – which can in turn be expressed just in terms of our poset: two separations are *nested* if they have comparable orientations.

⁴The most prominent of these, perhaps, will be tree-width and path-width.

⁵Note the similarity to ultrafilters, a standard kind of abstract HCSs in infinite contexts.

While our duality theorem for these abstract separation systems implies all the duality theorems mentioned so far, by applying it to separations in graphs or matroids, it can also be applied in very different contexts.

For example, the bipartitions of the set D of pixels of an image form a separation system: they are partially ordered by inclusion of their sides, and the involution of flipping the sides of the bipartition inverts this ordering. Depending on the application, some ways of cutting the image in two will be more natural than others, which gives rise to a cost function on these separations of D . Taking this cost of a separation as its ‘order’ then gives rise to tangles: abstract HCSs signifying regions of the image. Unlike regions defined by simply specifying a subset of D , regions defined by tangles are allowed to be fuzzy (as in our earlier grid example) – much like regions in real-world images.

If the cost function on the separations of our image is submodular – which in practice is not a severe restriction – our duality theorem can be applied to these tangles [6]. For every integer k , the application will either find a region of order at least k or produce a nested ‘tree’ set of bipartitions, all of order $< k$, which together witness that no such region exists [7]. This information could be used, for example, to assess the quality of an image, eg. after sending it through a noisy channel.

There are also potential applications in pure mathematics. For a very simple example, consider a triangulation of a topological sphere. This can be cut in two, in many ways, by closed paths along the edges of the triangulation, i.e., by cycles in its 1-skeleton. The lengths of these paths or cycles define a submodular order function on the separations of our sphere that they define. Our duality theorem then says that, for every integer k , there either exists a region dense enough that no closed path of length $< k$ can cut the sphere so as to divide this region roughly in half, or there exists a collection of non-crossing paths each of length $< k$ which, between them, cut up the entire sphere in a tree-like way (with a ternary tree) into single triangles. This can be done for other complexes too, but exploring this further would take us too far afield here.

Our abstract duality theorem will come in two flavours, ‘weak’ and ‘strong’. Our *weak duality theorem*, presented in Section 3, will be easy to prove but has no direct applications. It will be used as a stepping stone for the *strong duality theorem*, our main result, which we prove in Section 4. In Section 5 we present a refinement of the strong duality theorem.

2 Terminology and basic facts

A *separation of a set V* is a set $\{A, B\}$ such that $A \cup B = V$. The ordered pairs (A, B) and (B, A) are its *orientations*. The *oriented separations* of V are the orientations of its separations. Mapping every oriented separation (A, B) to its *inverse* (B, A) is an involution that reverses the partial ordering

$$(A, B) \leq (C, D) :\Leftrightarrow A \subseteq C \text{ and } B \supseteq D.$$

Note that this is equivalent to $(D, C) \leq (B, A)$. Informally, we think of (A, B) as *pointing towards B* and *away from A* . Similarly, if $(A, B) \leq (C, D)$, then (A, B) *points towards* $\{C, D\}$, while (C, D) *points away from* $\{A, B\}$.

Generalizing these properties of separations of sets, we now give an axiomatic definition of ‘abstract’ separations.⁶ A *separation system* $(\vec{S}, \leq, *)$ is a partially ordered set \vec{S} with an order-reversing involution $*$. Its elements are called *oriented separations*. When a given element of \vec{S} is denoted as \vec{s} , its *inverse* \vec{s}^* will be denoted as \bar{s} , and vice versa. The assumption that $*$ be *order-reversing* means that, for all $\vec{r}, \vec{s} \in \vec{S}$,

$$\vec{r} \leq \vec{s} \Leftrightarrow \bar{r} \geq \bar{s}. \quad (1)$$

A *separation* is a set of the form $\{\vec{s}, \bar{s}\}$, and then denoted by s . We call \vec{s} and \bar{s} the *orientations* of s . The set of all such sets $\{\vec{s}, \bar{s}\} \subseteq \vec{S}$ will be denoted by S . If $\vec{s} = \bar{s}$, we call both \vec{s} and s *degenerate*.

When a separation is introduced ahead of its elements and denoted by a single letter s , its elements will then be denoted as \vec{s} and \bar{s} .⁷ Given a set $S' \subseteq S$ of separations, we write $\vec{S}' := \bigcup S' \subseteq \vec{S}$ for the set of all the orientations of its elements. With the ordering and involution induced from \vec{S} , this is again a separation system.⁸

Separations of sets, and their orientations, are clearly an instance of this if we identify $\{A, B\}$ with $\{(A, B), (B, A)\}$.

If there are binary operations \vee and \wedge on our separation system \vec{S} such that $\vec{r} \vee \vec{s}$ is the supremum and $\vec{r} \wedge \vec{s}$ the infimum of \vec{r} and \vec{s} in \vec{S} , we call $(\vec{S}, \leq, *, \vee, \wedge)$ a *universe* of (oriented) separations. By (1), it satisfies De Morgan’s law:

$$(\vec{r} \vee \vec{s})^* = \bar{r} \wedge \bar{s}. \quad (2)$$

The oriented separations of a set V form such a universe: if $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$, say, then $\vec{r} \vee \vec{s} := (A \cup C, B \cap D)$ and $\vec{r} \wedge \vec{s} := (A \cap C, B \cup D)$ are again oriented separations of V , and are the supremum and infimum of \vec{r} and \vec{s} . Similarly, the oriented separations of a graph form a universe. Its oriented separations of order $< k$ for some fixed k , however, form a separation system inside this universe that may not itself be a universe with respect to \vee and \wedge as defined above.

A separation $\vec{r} \in \vec{S}$ is *trivial in \vec{S}* , and \vec{r} is *co-trivial*, if there exists $s \in S$ such that $\vec{r} < \vec{s}$ as well as $\vec{r} < \bar{s}$. Note that if \vec{r} is trivial in \vec{S} then so is every $\vec{r}' \leq \vec{r}$. If \vec{r} is trivial, witnessed by \vec{s} , then $\vec{r} < \vec{s} < \bar{r}$ by (1). Hence if \vec{r} is trivial, then \bar{r} cannot be trivial, because that would imply $\vec{r} < \bar{r} < \vec{r}$. In particular, trivial separations are never degenerate.

If \vec{S} is non-empty and finite, its separations cannot all be trivial or co-trivial: if \vec{r} is maximal among trivial ones, then both orientations of every $s \in S$ witnessing the triviality of \vec{r} are nontrivial.

There can also be nontrivial separations \vec{s} such that $\vec{s} < \bar{s}$. But anything smaller than these is again trivial: if $\vec{r} < \vec{s} \leq \bar{s}$, then s witnesses the triviality

⁶Our proofs will read a lot more smoothly in this abstract set-up, because it will use the familiar and intuitive arrow notation for orientations. When drawing examples, however, we shall usually work with separations of sets.

⁷It is meaningless here to ask which is which: neither \vec{s} nor \bar{s} is a well-defined object just given s . But given one of them, both the other and s will be well defined. They may be degenerate, in which case $s = \{\vec{s}\} = \{\bar{s}\}$.

⁸For $S' = S$, our definition of \vec{S}' is consistent with the existing meaning of \vec{S} . When we refer to oriented separations using explicit notation that indicates orientation, such as \vec{s} or (A, B) , we sometimes leave out the word ‘oriented’ to improve the flow of words. Thus, when we speak of a ‘separation (A, B) ’, this will in fact be an oriented separation.

of \vec{r} . Separations \vec{s} such that $\vec{s} \leq \vec{s}$, trivial or not, will be called *small*; note that, by (1), if \vec{s} is small then so is every $\vec{s}' \leq \vec{s}$. If \vec{r} is small but not trivial, then $\vec{r} < \vec{r}$ but there is no \vec{s} with $\vec{r} < \vec{s} < \vec{r}$, since any such s would make \vec{r} trivial.

The trivial oriented separations of a set V , for example, are those of the form $\vec{r} = (A, B)$ with $A \subseteq C \cap D$ and $B \supseteq C \cup D = V$ for some $s = \{C, D\} \neq r$. The small separations (A, B) of V are all those with $B = V$.

Two separations r, s are *nested* if they have comparable orientations; otherwise they *cross*. Two oriented separations \vec{r}, \vec{s} are *nested* if r and s are nested.⁹ We say that \vec{r} *points towards* s , and \vec{r} *points away from* s , if $\vec{r} \leq \vec{s}$ or $\vec{r} \leq \vec{s}$. Then two nested oriented separations are either comparable, or point towards each other, or point away from each other. A set of separations is *nested* if every two of its elements are nested.

A set $O \subseteq \vec{S}$ of oriented separations is *antisymmetric* if it does not contain the inverse of any of its nondegenerate elements. It is *consistent* if there are no distinct $r, s \in S$ with orientations $\vec{r} < \vec{s}$ such that $\vec{r}, \vec{s} \in O$. (Informally: if it does not contain orientations of distinct separations that point away from each other.)

A set σ of nondegenerate oriented separations, possibly empty, is a *star of separations* if they point towards each other: if $\vec{r} \leq \vec{s}$ for all distinct $\vec{r}, \vec{s} \in \sigma$ (Fig. 1). Stars of separations are clearly nested. They are also consistent: if \vec{r}, \vec{s} lie in the same star we cannot have $\vec{r} < \vec{s}$, since also $\vec{s} \leq \vec{r}$ by the star property. A star σ need not be antisymmetric; but if $\{\vec{s}, \vec{s}\} \subseteq \sigma$, then any other $\vec{r} \in \sigma$ will be trivial.




Figure 1: The separations $(A, B), (C, D), (E, F)$ form a 3-star

An *orientation* of a set S of separations is a set $O \subseteq \vec{S}$ that contains for every $s \in S$ exactly one of its orientations \vec{s}, \vec{s} . A *partial orientation* of S is an orientation of a subset of S , i.e., an antisymmetric subset of \vec{S} .

Every consistent orientation of S contains all separations \vec{r} that are trivial in \vec{S} , because it cannot contain their inverse \vec{r} : if the triviality of \vec{r} is witnessed by $s \in S$, say, then \vec{r} would be inconsistent with both \vec{s} and \vec{s} . It is not hard to show that every consistent partial orientation of S containing no co-trivial $\vec{r} \in \vec{S}$ extends to a consistent orientation of all of S [3].

⁹Terms introduced for unoriented separations may be used informally for oriented separations too if the meaning is obvious, and vice versa.

Let S be a set of separations. An S -tree is a pair (T, α) of a tree¹⁰ T and a function $\alpha: \vec{E}(T) \rightarrow \vec{S}$ from the set

$$\vec{E}(T) := \{ (x, y) : \{x, y\} \in E(T) \}$$

of the *orientations* (x, y) of its edges $\{x, y\}$ to \vec{S} such that, for every edge e of T , if α maps $\vec{e} = (x, y)$ to \vec{s} it maps its inverse $\vec{e} := (y, x)$ to \vec{s} . It is an S -tree is over $\mathcal{F} \subseteq 2^{\vec{S}}$ if, in addition, for every node t of T we have $\alpha(\vec{F}_t) \in \mathcal{F}$, where

$$\vec{F}_t := \{ (x, t) : xt \in E(T) \}.$$

We shall call the set $\vec{F}_t \subseteq \vec{E}(T)$ the *oriented star at t* in T . Its image $\alpha(\vec{F}_t) \in \mathcal{F}$ is said to be *associated with t* in (T, α) .

The S -tree (T, α) is *redundant* if it has a node t of T with distinct neighbours t', t'' such that $\alpha(t', t) = \alpha(t'', t)$; otherwise it is *irredundant*. Redundant S -trees can clearly be *pruned* to irredundant ones over the same \mathcal{F} :

Lemma 2.1. [3] *If (T, α) is an S -tree over \mathcal{F} , possibly redundant, and x is any node of T , then T has a subtree T' containing x such that (T', α') is an irredundant S -tree over \mathcal{F} , where α' is the restriction of α to $\vec{E}(T')$. \square*

An important example of S -trees are irredundant S -trees *over stars*: those over some \mathcal{F} all of whose elements are stars of separations. In such an S -tree (T, α) the map α preserves the *natural partial ordering* on $\vec{E}(T)$ defined by letting $(x, y) < (u, v)$ if $\{x, y\} \neq \{u, v\}$ and the unique $\{x, y\}$ - $\{u, v\}$ path in T joins y to u (see Figure 2):




Figure 2: Edges $(x, y) < (u, v)$ and separations $(A, B) = \alpha(x, y) \leq \alpha(u, v) = (C, D)$

Lemma 2.2. *Let (T, α) be an irredundant S -tree over a set \mathcal{F} of stars of separations. Then α preserves the natural partial ordering on $\vec{E}(T)$.*

Proof. As (T, α) is irredundant, distinct edges in an oriented star \vec{F}_t in T map to distinct separations in $\alpha(\vec{F}_t)$. These point towards each other, since $\alpha(\vec{F}_t) \in \mathcal{F}$ is a star. Formally, this means that for every $t \in T$ the map α preserves the partial ordering \leq that $\vec{E}(T)$ induces on the elements of \vec{F}_t and their inverses. This propagates through $\vec{E}(T)$, to the effect that α preserves this ordering on all of $\vec{E}(T)$: whenever $\vec{e} \leq \vec{f}$ in $\vec{E}(T)$ we have $\alpha(\vec{e}) \leq \alpha(\vec{f})$ in \vec{S} . \square

¹⁰Trees have at least one node [4].

Two edges of an irredundant S -tree over stars cannot have orientations that point towards each other and map to the same separation, unless this is trivial:

Lemma 2.3. *Let (T, α) be an irredundant S -tree over a set \mathcal{F} of stars. Let e, f be distinct edges of T with orientations $\vec{e} < \vec{f}$ such that $\alpha(\vec{e}) = \alpha(\vec{f}) =: \vec{r}$. Then \vec{r} is trivial.*

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

Proof. If α maps all \vec{e}' with $\vec{e} < \vec{e}' < \vec{f}$ to \vec{r} or to \vec{r} , then the e - f path in T has a node with two incoming edges mapped to \vec{r} . This contradicts our assumption that (T, α) is irredundant. Hence there exists such an edge \vec{e}' with $\alpha(\vec{e}') = \vec{s}$ for some $s \neq r$. Lemma 2.2 implies that $\vec{r} = \alpha(\vec{e}) \leq \alpha(\vec{e}') \leq \alpha(\vec{f}) = \vec{r}$, so $\vec{r} \leq \vec{s}$ as well as $\vec{r} \leq \vec{s}$ by (1). As $s \neq r$ these inequalities are strict, so s witnesses that \vec{r} is trivial. \square

Recall that stars of separations need not, by definition, be antisymmetric. While it is important for our proofs to allow this, we can always contract an S -tree (T, α) over a set \mathcal{F} of stars to an S -tree (T', α') over the subset $\mathcal{F}' \subseteq \mathcal{F}$ of its antisymmetric stars. Indeed, if T has a node t such that $\alpha(\vec{F}_t)$ is not antisymmetric, then t has neighbours t', t'' such that $\alpha(t', t) = \vec{s} = \alpha(t, t'')$ for some $\vec{s} \in \vec{S}$. Let T' be the tree obtained from T by deleting the component of $T - t't - tt''$ containing t and joining t' to t'' . Let $\alpha'(t', t'') := \vec{s}$ and $\alpha'(t'', t') := \vec{s}$, and otherwise let $\alpha' := \alpha \upharpoonright \vec{E}(T')$. Then (T', α') is again an S -tree over \mathcal{F} . Since we can do this whenever some \vec{F}_t maps to a star of separations that is not antisymmetric, but only finitely often, we must arrive at an S -tree over \mathcal{F}' .

Lemma 2.4. *Let (T, α) be an S -tree over a set \mathcal{F} of stars. Let x be a leaf of T , let e be the edge of T at x , and let \vec{e} be oriented away from x . Assume that $\vec{r} = \alpha(\vec{e})$ is nontrivial. Then T has a minor T' containing x such that (T', α') , where $\alpha' = \alpha \upharpoonright \vec{E}(T')$, is an irredundant S -tree over \mathcal{F} , and such that \vec{e} is the only edge in $\vec{E}(T')$ with $\alpha'(\vec{e}) = \vec{r}$.*

Proof. By Lemma 2.1 we may assume that (T, α) is irredundant. Suppose T has an edge f with an orientation $\vec{f} \neq \vec{e}$ such that $\alpha(\vec{f}) = \vec{r}$. Then $f \neq e$, since otherwise $\vec{f} = \vec{e}$ and hence $\alpha(\vec{f}) = \vec{r}$ as well as $\alpha(f) = \vec{r}$, which would make r degenerate and thus contradict our assumption that $\{\vec{r}\} \in \mathcal{F}$ is a star.

By Lemma 2.3 we cannot have $\vec{e} < \vec{f}$, so $\vec{e} < \vec{f}$ since \vec{e} issues from a leaf. By Lemma 2.2, every edge \vec{e}' with $\vec{e} \leq \vec{e}' \leq \vec{f}$ satisfies $\vec{r} = \alpha(\vec{e}) \leq \alpha(\vec{e}') \leq \alpha(\vec{f}) = \vec{r}$, so $\alpha(\vec{e}') = \vec{r}$. We can now apply the reduction described just before this lemma to the initial node t of \vec{f} , to obtain a smaller S -tree over \mathcal{F} that is still irredundant and contains x and e . Iterating, we obtain the desired S -tree (T', α') . \square

3 Weak duality

Our paradigm in this paper is to capture the notion of ‘highly connected substructures’ in a given combinatorial structure by orientations O of a set S of separations of this structure that satisfy certain consistency rules laid down by specifying a set \mathcal{F} of ‘forbidden’ sets of oriented separations that O must not contain.

Let us say that a partial orientation P of S avoids $\mathcal{F} \subseteq 2^{\vec{S}}$ if $2^P \cap \mathcal{F} = \emptyset$.

Theorem 3.1 (Weak Duality Theorem). *Let $(\vec{S}, \leq, *)$ be a separation system and $\mathcal{F} \subseteq 2^{\vec{S}}$ a set of stars. Then exactly one of the following assertions holds:*

- (i) *There exists an S -tree over \mathcal{F} .*
- (ii) *There exists an orientation of S that avoids \mathcal{F} .*

The proof of the following lemma uses the simple fact that every orientation of a finite tree has a sink. To find one, just follow a directed path.

Lemma 3.2. *Let $(\vec{S}, \leq, *)$ be a separation system and $\mathcal{F} \subseteq 2^{\vec{S}}$. If there exists an S -tree over \mathcal{F} , then no orientation of S avoids \mathcal{F} .*

Proof. Let (T, α) be an S -tree over \mathcal{F} , and let O be an orientation of S . Let $t \in V(T)$ be a sink in the orientation of the edges of T that O induces via α . Then $\alpha(\vec{F}_t) \subseteq O$. Since $\alpha(\vec{F}_t) \in \mathcal{F}$, as (T, α) is an S -tree over \mathcal{F} , this means that O does not avoid \mathcal{F} . \square

Proof of Theorem 3.1. By Lemma 3.2, at most one of (i) and (ii) holds. We now show that at least one of them holds. Let

$$O^- := \{ \vec{s} \mid \{ \vec{s} \} \in \mathcal{F} \}.$$

Then any \mathcal{F} -avoiding orientation of S must include O^- as a subset. As \mathcal{F} consists of stars, O^- contains no degenerate separations.

If $O^- \supseteq \{ \vec{s}, \bar{s} \}$ for some $s \in S$, then (T, α) with $T = K_2$ and $\text{im } \alpha = \{ \vec{s}, \bar{s} \}$ is an S -tree over \mathcal{F} . So we may assume that O^- is antisymmetric: a partial orientation of $S \setminus D$, where D is the set of degenerate elements of S . We apply induction on $|S \setminus D| - |O^-|$ to show that, whenever \mathcal{F} is such that O^- is antisymmetric, either (i) or (ii) holds.

If $|S \setminus D| = |O^-|$, then $O^- \cup \vec{D}$ is an orientation of all of S . If (ii) fails then $O^- \cup \vec{D}$ has a subset $\sigma \in \mathcal{F}$. As \mathcal{F} consists of stars we have $\sigma \cap \vec{D} = \emptyset$, so $\sigma \subseteq O^-$. By definition of O^- , and since O^- is antisymmetric, σ is not a singleton set (though it may be empty). Let T be a star of $|\sigma|$ edges with centre t , say, and let α map its oriented edges (x, t) bijectively to σ . Then (T, α) satisfies (i).

We may thus assume that $S \setminus D$ contains a separation s_0 such that neither \vec{s}_0 nor \bar{s}_0 is in O^- . Let

$$\mathcal{F}_1 := \mathcal{F} \cup \{ \{ \vec{s}_0 \} \} \quad \text{and} \quad \mathcal{F}_2 := \mathcal{F} \cup \{ \{ \bar{s}_0 \} \},$$

and put $O_i^- := \{ \vec{s} \mid \{ \vec{s} \} \in \mathcal{F}_i \}$ for $i = 1, 2$. Note that $|O_i^-| > |O^-|$, and O_i^- is another partial orientation of $S \setminus D$.

Since any \mathcal{F}_i -avoiding orientation of S also avoids \mathcal{F} , we may assume that no orientation of S avoids \mathcal{F}_i , for both $i = 1, 2$. By the induction hypothesis, there are S -trees (T_i, α_i) over \mathcal{F}_i . Unless one of these is in fact over \mathcal{F} , the tree T_1 has a leaf x_1 associated with $\{ \vec{s}_0 \}$, while T_2 has a leaf x_2 associated with $\{ \bar{s}_0 \}$. Use Lemma 2.1 to prune the (T_i, α_i) to irredundant S -trees (T'_i, α'_i) over \mathcal{F}_i containing x_i . Suppose first that s_0 is nontrivial. Then Lemma 2.3 implies that x_1 and x_2 are the only leaves of T'_1 and T'_2 associated with $\{ \vec{s}_0 \}$ and $\{ \bar{s}_0 \}$, respectively.

Let T be the tree obtained from the disjoint union of $T'_1 - x_1$ and $T'_2 - x_2$ by joining the neighbour y_1 of x_1 in T'_1 to the neighbour y_2 of x_2 in T'_2 . Let

$\alpha: \vec{E}(T) \rightarrow \vec{S}$ map (y_1, y_2) to \vec{s}_0 and otherwise extend α'_1 and α'_2 . Then $\alpha'_1(y_1, x_1) = \alpha(y_1, y_2) = \alpha'_2(x_2, y_2)$, so α maps the oriented stars of edges at y_1 and y_2 to the same stars of separations in \vec{S} as α'_1 and α'_2 did. These lie in \mathcal{F} , so (T, α) is an S -tree over \mathcal{F} .

Suppose now that \vec{s}_0 , say, is trivial. Then \vec{s}_0 is non-trivial, and x_1 is the only leaf of T'_1 associated with $\{\vec{s}_0\}$, by Lemma 2.3 as before. Let x_1^1, \dots, x_1^n be the leaves of T'_2 associated with $\{\vec{s}_0\}$, and let T be obtained from the union of $T'_2 - \{x_1^1, \dots, x_1^n\}$ with n copies of $T'_1 - x_1$ by joining, for all $i = 1, \dots, n$, the neighbour y_1^i of x_1 in the i th copy of $T'_1 - x_1$ to the neighbour y_2^i of x_1^i in T'_2 . Define $\alpha: \vec{E}(T) \rightarrow \vec{S}$ as earlier, mapping (y_1^i, y_2^i) to \vec{s}_0 and (y_2^i, y_1^i) to \vec{s}_0 for all i , and otherwise extending α'_1 and α'_2 . \square

4 Strong duality

Theorem 3.1, alas, has a serious shortcoming: there are few, if any, sets S and $\mathcal{F} \subseteq 2^{\vec{S}}$ such that \mathcal{F} consists of stars in \vec{S} and the \mathcal{F} -avoiding orientations of S (all of them) capture an interesting notion of highly connected substructure found in the wild. The reason for this is that we are not, so far, requiring these orientations O to be consistent: we allow that O contains separations \vec{r} and \vec{s} when $\vec{r} < \vec{s}$, which will not usually be the case when O is induced by a meaningful highly connected substructure in the way discussed earlier. (We cannot simply add such sets $\{\vec{r}, \vec{s}\}$ to \mathcal{F} , since in order to be able to use Lemma 2.3 we must assume that \mathcal{F} consists of stars of separations.)

So what happens if we strengthen (ii) so as to ask for a consistent orientation of S ? Let us call a consistent \mathcal{F} -avoiding orientation of S an \mathcal{F} -tangle. Since all consistent orientations of S will contain all trivial $\vec{s} \in \vec{S}$, we may then add all co-trivial singletons $\{\vec{s}\}$ to \mathcal{F} without impeding the existence of an \mathcal{F} -tangle; this might help us find an S -tree over \mathcal{F} if no such orientation exists.

Still, our proof breaks down as early as the induction start: we now also have to ask that O^- – indeed, $O^- \cup \vec{D}$ – should be consistent. It is not even unnatural to ensure that O^- is closed down in (\vec{S}, \leq) (which implies consistency), by requiring that if $\{\vec{r}\} \in \mathcal{F}$ and $\vec{r} < \vec{s}$ then also $\{\vec{s}\} \in \mathcal{F}$. For if a singleton star $\{\vec{r}\}$ is in \mathcal{F} , the idea is that the part of our structure to which \vec{r} points is too small to contain a highly connected substructure; and then the same should apply to all $\vec{s} \geq \vec{r}$.

But now we have a problem at the induction step: when forming the \mathcal{F}_i , we now have to add not only $\{\vec{s}_0\}$ and $\{\vec{s}_0\}$ to \mathcal{F} , but all singleton stars $\{\vec{s}\}$ with $\vec{s}_0 \leq \vec{s}$ or $\vec{s}_0 \leq \vec{s}$, respectively, to keep O^- closed down. This, then, spawns more problems: now both T_i can have many leaves associated with a singleton star of \mathcal{F}_i that is not in \mathcal{F} . Even if each of these occurs at most once, there is no longer an obvious way of how to merge T_1 and T_2 into a single S -tree over \mathcal{F} .

We shall deal with this problem as follows. Rather than adding singletons of the orientations of some fixed separation s_0 to \mathcal{F} to form the \mathcal{F}_i , we shall add some $\{\vec{r}_i\}$ such that \vec{r}_i is minimal in $\vec{S} \setminus (O^- \cup \vec{D})$. Then $\{\vec{r}_i\}$ can be associated with at most one leaf of (T'_i, α'_i) , because $\vec{r}_i \notin O^-$ will be nontrivial (cf. Lemma 2.3). These \vec{r}_i will be chosen nested, so that both point to some s_0 between them. We shall then modify the two S -trees over the \mathcal{F}_i into S -trees over $\mathcal{F} \cup \{\{\vec{s}_0\}\}$ and $\mathcal{F} \cup \{\{\vec{s}_0\}\}$ by ‘shifting’ their separations to either side of s_0 , and then merge these shifted S -trees as before to obtain one over \mathcal{F} .

Let \bar{r} be a nontrivial and nondegenerate element of a separation system $(\bar{S}, \leq, *)$ contained in some universe $(\bar{U}, \leq, *, \vee, \wedge)$ of separations, the ordering and involution on \bar{S} being induced by those of \bar{U} . Later,¹¹ $\{\bar{r}\}$ will be associated with a leaf of an S -tree (T, α) over some set $\mathcal{F} \subseteq 2^{\bar{S}}$ of stars. In particular, $\{\bar{r}\}$ will be a star, so r will be nondegenerate.

Consider any $\bar{s}_0 \in \bar{S}$ such that $\bar{r} \leq \bar{s}_0$. As \bar{r} is nontrivial and nondegenerate, so is \bar{s}_0 . Our aim will be to ‘shift’ (T, α) to a new S -tree (T, α') based on the same tree T , by modifying α in such a way that \bar{s}_0 points to all the separations in the image of α' .

Let $S_{\geq \bar{r}}$ be the set of all separations $s \in S$ that have an orientation $\bar{s} \geq \bar{r}$. Since $\{\bar{r}\}$ will be associated with a leaf of T , and the T -edge at that leaf, when oriented away from it, points to all the other edges of T (Fig. 3, left), Lemma 2.2 implies that $S_{\geq \bar{r}}$ will include the image of α . Since \bar{r} is nontrivial, only one of the two orientations \bar{s} of every $s \in S_{\geq \bar{r}} \setminus \{\bar{r}\}$ satisfies $\bar{s} \geq \bar{r}$. Letting

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

for all $\bar{s} \geq \bar{r}$ in $S_{\geq \bar{r}} \setminus \{\bar{r}\}$ thus defines a map $S_{\geq \bar{r}} \rightarrow \bar{U}$, the *shifting map* $f \downarrow_{\bar{s}_0}^{\bar{r}}$ (Fig. 3, right). Note that $f \downarrow_{\bar{s}_0}^{\bar{r}}(\bar{r}) = \bar{s}_0$, since $\bar{r} \leq \bar{s}_0$.




Figure 3: Shifting $\alpha(\bar{e}) = \bar{s}$ to $\alpha'(\bar{e}) = \bar{s} \vee \bar{s}_0$

The concatenation

$$\alpha' := f \downarrow_{\bar{s}_0}^{\bar{r}} \circ \alpha$$

will thus be a well-defined map from $\bar{E}(T)$ to \bar{U} . We would like its images to lie in \bar{S} , but will need some assumptions to ensure this.

To help us show that (T, α') is over \mathcal{F} if (T, α) is, we also need that $f \downarrow_{\bar{s}_0}^{\bar{r}}$ maps the stars $\alpha(\bar{F}_t)$ to stars. This would be immediate if we could show that $f \downarrow_{\bar{s}_0}^{\bar{r}}$ preserved the ordering \leq on $\bar{E}(T)$. It does in fact do this, but with one possible exception: if $\bar{r} < \bar{r}$, which can happen, we shall usually have $f \downarrow_{\bar{s}_0}^{\bar{r}}(\bar{r}) = \bar{s}_0 \not\leq \bar{s}_0 = f \downarrow_{\bar{s}_0}^{\bar{r}}(\bar{r})$.¹² By choosing (T, α) irredundant we shall be able to ensure that this exception remains irrelevant, but until then we have to exclude it:

Lemma 4.1. *The map $f = f \downarrow_{\bar{s}_0}^{\bar{r}}$ preserves the ordering \leq on $S_{\geq \bar{r}} \setminus \{\bar{r}\}$. In particular, f maps stars to stars.*

¹¹To motivate our definitions, we shall keep looking ahead to their intended application, which will involve an S -tree (T, α) . But formally we only need the assumptions stated earlier: that \bar{r} is nontrivial and nondegenerate.

¹²This is because $f \downarrow_{\bar{s}_0}^{\bar{r}}(\bar{r})$, unlike $f \downarrow_{\bar{s}_0}^{\bar{r}}(\bar{s})$ for all other $\bar{s} \geq \bar{r}$, was defined indirectly as the inverse of $f \downarrow_{\bar{s}_0}^{\bar{r}}(\bar{r})$. Note that Lemma 4.1 excludes *only* this case: every $\bar{s} \in S_{\geq \bar{r}}$ other than \bar{r} itself relates to \bar{r} by its relationship to \bar{r} via (1), which $f \downarrow_{\bar{s}_0}^{\bar{r}}$ preserves.

Proof. Consider separations $\vec{s} \leq \vec{s}'$ in $\vec{S}_{\geq \vec{r}} \setminus \{\vec{r}\}$. Suppose first that $\vec{r} \leq \vec{s}$. Then $f(\vec{s}') = \vec{s}' \vee \vec{s}_0$ lies above (is \geq than) both \vec{s}_0 and \vec{s} . It therefore also lies above their supremum $f(\vec{s}) = \vec{s} \vee \vec{s}_0$. We may therefore assume that $\vec{r} \not\leq \vec{s}$.

If $\vec{r} < \vec{s}'$ then $\vec{r} < \vec{s}' \leq \vec{s}$, which reduces to the case above on renaming \vec{s}' as \vec{s} and \vec{s} as \vec{s}' . We may thus assume that $\vec{r} \not\leq \vec{s}'$, and hence $\vec{r} \leq \vec{s}'$.

Now $f(\vec{s}) = (\vec{s} \vee \vec{s}_0)^* = \vec{s} \wedge \vec{s}_0 \leq \vec{s} \leq \vec{s}' \leq \vec{s}' \vee \vec{s}_0 = f(\vec{s}')$ by (2).

For a proof of our claim that f maps stars (including $\{\vec{r}\}$) to stars, we still need to check that no separation in the image of f is degenerate. But every $s \in S_{\geq \vec{r}}$ has an orientation \vec{s} such that $\vec{r} \leq \vec{s} \leq f(\vec{s})$. Since \vec{r} is nontrivial and nondegenerate, this makes $f(\vec{s})$ nondegenerate. (Indeed, if $f(\vec{s}) = \vec{t}$ is degenerate then $r \neq t$, since r is nondegenerate. But then $\vec{r} < \vec{t} = \vec{t}$, making \vec{r} trivial.) \square

Let us say that $\vec{s}_0 \in \vec{S}$ *emulates* $\vec{r} \in \vec{S}$ in \vec{S} if $\vec{s}_0 \geq \vec{r}$ and every $\vec{s} \in \vec{S} \setminus \{\vec{r}\}$ with $\vec{s} \geq \vec{r}$ satisfies $\vec{s} \vee \vec{s}_0 \in \vec{S}$. (This will ensure that α' has its image in \vec{S} .) Let us call \vec{S} *separable* if for every two nontrivial and nondegenerate $\vec{r}, \vec{r}' \in \vec{S}$ such that $\vec{r} \leq \vec{r}'$ there exists an $s_0 \in S$ with an orientation \vec{s}_0 emulating \vec{r} in \vec{S} and its inverse \vec{s}_0 emulating \vec{r}' in \vec{S} . (Notice that any such s_0 will also be nontrivial and nondegenerate.¹³)

Finally, we need a condition on \mathcal{F} to ensure that the shifts of stars of separations associated with nodes of T are not only again stars but are also again in \mathcal{F} . Given any set $\mathcal{F} \subseteq 2^{\vec{U}}$ of stars, let us say that a separation $\vec{s}_0 \in \vec{S}$ *emulates* $\vec{r} \in \vec{S}$ in \vec{S} for \mathcal{F} if \vec{s}_0 emulates \vec{r} in \vec{S} and for any star $\sigma \subseteq \vec{S}_{\geq \vec{r}} \setminus \{\vec{r}\}$ in \mathcal{F} that has an element $\vec{s} \geq \vec{r}$ we also have $f \downarrow_{\vec{s}_0}^{\vec{r}}(\sigma) \in \mathcal{F}$.¹⁴

Let us say that \mathcal{F} *forces* the separations $\vec{s} \in \vec{S}$ for which $\{\vec{s}\} \in \mathcal{F}$. And that \vec{S} is \mathcal{F} -*separable* if for all nontrivial and nondegenerate $\vec{r}, \vec{r}' \in \vec{S}$ that are not forced by \mathcal{F} and satisfy $\vec{r} \leq \vec{r}'$ there exists an $s_0 \in S$ with an orientation \vec{s}_0 that emulates \vec{r} in \vec{S} for \mathcal{F} and such that \vec{s}_0 emulates \vec{r}' in \vec{S} for \mathcal{F} . (As earlier, any such \vec{s}_0 will also be nontrivial and nondegenerate.)

We now have all the ingredients needed to shift an S -tree:

Lemma 4.2. *Let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a set of stars, and let (T, α) be an irredundant S -tree over \mathcal{F} . Let x be a leaf of T , and let \vec{e} be its edge at x oriented away from x . Assume that $\vec{r} = \alpha(\vec{e})$ is nontrivial, and that $\alpha(\vec{e}') \neq \vec{r}$ for every other $\vec{e}' \in \vec{E}(T)$. Let $\vec{s}_0 \in \vec{S}$ emulate \vec{r} in \vec{S} for \mathcal{F} , and let $\alpha' := f \downarrow_{\vec{s}_0}^{\vec{r}} \circ \alpha$. Then (T, α') is an S -tree over $\mathcal{F} \cup \{\{\vec{s}_0\}\}$ in which $\{\vec{s}_0\}$ is a star associated with x but with no other leaf of T .*

Proof. Since \mathcal{F} consists of stars and (T, α) is irredundant, the map α preserves the natural ordering on $\vec{E}(T)$ (Lemma 2.2). Since $\{\vec{r}\}$ is associated with a leaf of T , and every edge of T has an orientation away from this leaf, α maps $\vec{E}(T)$ to $\vec{S}_{\geq \vec{r}}$. As $\{\vec{r}\}$ is a star, r is nondegenerate, so $f \downarrow_{\vec{s}_0}^{\vec{r}}$ is correctly defined on $\vec{S}_{\geq \vec{r}}$ and α' is well defined. Its image lies in \vec{S} , because \vec{s}_0 emulates \vec{r} .

By Lemma 4.1, $f \downarrow_{\vec{s}_0}^{\vec{r}}$ maps stars to stars. As x is the only node t of T with

¹³However it can happen that $\vec{r} < \vec{r}' \leq \vec{s}_0$ ($\leq \vec{r}'$). Then $\vec{r}' \leq \vec{s}_0 \leq \vec{r}$ as well as, by assumption, $\vec{r}' \leq \vec{r}$, so the nontriviality of \vec{r}' implies that $r = r'$. Then $\vec{r} < \vec{r}' \leq \vec{s}_0 \leq \vec{r}'$ with equality in both cases, giving $\vec{r} = \vec{r}'$.

¹⁴In fact, we could make do with less: that $f \downarrow_{\vec{s}_0}^{\vec{r}}$ is defined (with image in \vec{S} , as now) only on some symmetric subset of $\vec{S}_{\geq \vec{r}}$ that contains $\bigcup \mathcal{F}$, and that for some fixed $S' \subseteq S$ and every σ as above we have $f \downarrow_{\vec{s}_0}^{\vec{r}}(\sigma) \cap \vec{S}' \in \mathcal{F}$ if $f \downarrow_{\vec{s}_0}^{\vec{r}}(\vec{s}) \in \vec{S}'$ for the unique $\vec{s} \geq \vec{r}$ in σ . It will be easy to adapt the proof of Theorem 4.3 should this ever be necessary. See Section 5 for more on this.

$\vec{r} \in \alpha(\vec{F}_t)$ and s_0 emulates \vec{r} for \mathcal{F} with $t \neq x$ to stars in \mathcal{F} . As x is associated with $\{\vec{r}\}$ in (T, α) and $f \downarrow_{\vec{s}_0}^{\vec{r}}(\vec{r}) = \vec{s}_0$, this means that (T, α') is an S -tree over $\mathcal{F} \cup \{\{\vec{s}_0\}\}$ in which x is associated with $\{\vec{s}_0\}$ (which is also a star).

Suppose $\{\vec{s}_0\}$ is also associated in (T, α') with another leaf $y \neq x$ of T . Let $\{\vec{r}'\}$ be associated with y in (T, α) . Then $f \downarrow_{\vec{s}_0}^{\vec{r}'}(\vec{r}') = \vec{s}_0$, while $\vec{r}' \leq \vec{r}$ by Lemma 2.2. If $\vec{r} = \vec{s}_0$ then $f \downarrow_{\vec{s}_0}^{\vec{r}}$ is the identity on $\vec{S}_{\geq \vec{r}}$, so $\{\vec{s}_0\} = \{\vec{r}\}$ is associated with x and y also in (T, α) . But this contradicts our assumptions about \vec{r} . Hence $\vec{r} < \vec{s}_0 = f \downarrow_{\vec{s}_0}^{\vec{r}'}(\vec{r}') \leq f \downarrow_{\vec{s}_0}^{\vec{r}}(\vec{r}) = \vec{s}_0$, where the \leq comes from applying Lemma 4.1 to $\vec{r} \leq \vec{r}'$ (neither of which is \vec{r} , by assumption) and then applying (1). Thus, s_0 witnesses that \vec{r} is trivial in \vec{S} , contrary to our assumptions. \square

We can now strengthen our weak duality theorem so as to yield consistent orientations, provided that \vec{S} is \mathcal{F} -separable. Recall that for a separation system $(\vec{S}, \leq, *)$ and a set \mathcal{F} , an orientation O of S is called an \mathcal{F} -tangle if it is consistent and avoids \mathcal{F} , that is, if $2^O \cap \mathcal{F} = \emptyset$. Let us call \mathcal{F} *standard* for \vec{S} if it forces all $\vec{s} \in \vec{S}$ that are trivial in \vec{S} .

Theorem 4.3 (Strong Duality Theorem). *Let $(\vec{U}, \leq, *, \vee, \wedge)$ be a universe of separations containing a separation system $(\vec{S}, \leq, *)$. Let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars, standard for \vec{S} . If \vec{S} is \mathcal{F} -separable, exactly one of the following assertions holds:*

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

Proof. Since replacing \mathcal{F} with $\mathcal{F} \cap 2^{\vec{S}}$ leaves both (i) and (ii) unchanged we may, and shall, assume that $\mathcal{F} \subseteq 2^{\vec{S}}$. By Lemma 3.2, (i) and (ii) cannot both hold; we show that at least one of them holds.

Since \mathcal{F} is standard, the set

$$O^- := \{ \vec{s} \mid \{\vec{s}\} \in \mathcal{F} \}$$

of separations that \mathcal{F} forces contains all the trivial separations in \vec{S} . But it contains no degenerate ones, because the $\{\vec{s}\} \in \mathcal{F}$ are stars. If $O^- \supseteq \{\vec{s}, \vec{s}\}$ for some $s \in S$, then (T, α) with $T = K_2$ and $\text{im } \alpha = \{\vec{s}, \vec{s}\}$ is an S -tree over \mathcal{F} . We may thus assume that O^- is antisymmetric: a partial orientation of $S \setminus D$, where D is the set of degenerate elements of S .

Let us show that O^- is consistent. If not, then O^- contains some \vec{r} and \vec{r}' such that $\vec{r} < \vec{r}'$. As O^- is antisymmetric, it then does not contain their inverses \vec{r} and \vec{r}' . So \mathcal{F} does not force these; in particular, they are nontrivial. Since \vec{S} is \mathcal{F} -separable, there exists an $s_0 \in S$ with orientations \vec{s}_0, \vec{s}_0 such that \vec{s}_0 emulates \vec{r} in \vec{S} for \mathcal{F} and \vec{s}_0 emulates \vec{r}' in \vec{S} for \mathcal{F} . Since r is not degenerate and \vec{s}_0 emulates \vec{r} for \mathcal{F} , the singleton star $\{\vec{r}\} \in \mathcal{F}$ shifts to $\{f \downarrow_{\vec{s}_0}^{\vec{r}}(\vec{r})\} = \{\vec{s}_0\} \in \mathcal{F}$, so $\vec{s}_0 \in O^-$. Likewise, since r' is not degenerate and \vec{s}_0 emulates \vec{r}' we have $\vec{s}_0 \in O^-$. This contradicts our assumption that O^- is antisymmetric.

Let us show that $O^- \cup \vec{D}$ is still consistent. Suppose $r \neq s$ are such that $\vec{r}, \vec{s} \in O^- \cup \vec{D}$ and $\vec{r} < \vec{s}$. Then r and s are not both in D , since that would imply $\vec{r} < \vec{s} = \vec{s} < \vec{r} = \vec{r}$. Since O^- is consistent, we may thus assume that $\vec{r} \in O^-$ and $\vec{s} \in \vec{D}$ (or vice versa, which is equivalent by (1)). Then \vec{r} is trivial, as $\vec{r} < \vec{s} = \vec{s}$. Hence $\vec{r} \in O^-$ as well as, by assumption, $\vec{r} \in O^-$. This contradicts our assumption that O^- is antisymmetric.

Let R be the set of separations in $S \setminus D$ of which neither orientation lies in O^- . We shall apply induction on $|R|$ to show that (i) or (ii) holds whenever \mathcal{F} is such that O^- is antisymmetric. If $|R| = 0$, then $O^- \cup \vec{D}$ is either itself an \mathcal{F} -tangle of S or contains a star $\sigma \in \mathcal{F}$. Then $\sigma \subseteq O^-$, since stars have no degenerate elements. By definition of O^- , and since O^- is antisymmetric, σ is not a singleton subset of O^- (though it may be empty). Let T be a star of $|\sigma|$ edges with centre t , say, and let α map its oriented edges (x, t) bijectively to σ . Then (T, α) satisfies (ii).

For the induction step, pick $\vec{r}_0 \in \vec{R}$. Then neither \vec{r}_0 nor \vec{r}_0 lies in $O^- \cup \vec{D}$; let $\vec{r}_1 \leq \vec{r}_0$ and $\vec{r}_2 \leq \vec{r}_0$ be minimal in $\vec{S} \setminus (O^- \cup \vec{D})$. Then $\vec{r}_1 \leq \vec{r}_0 \leq \vec{r}_2$. As \mathcal{F} forces neither \vec{r}_1 nor \vec{r}_2 and \vec{S} is \mathcal{F} -separable, there exists an $s_0 \in S \setminus D$ with nontrivial orientations \vec{s}_0, \bar{s}_0 such that \vec{s}_0 emulates \vec{r}_1 in \vec{S} for \mathcal{F} and \bar{s}_0 emulates \vec{r}_2 in \vec{S} for \mathcal{F} .¹⁵

Since O^- is antisymmetric, it does not contain both \vec{s}_0 and \bar{s}_0 . Let us assume that $\bar{s}_0 \notin O^-$, i.e. that $\{\bar{s}_0\} \notin \mathcal{F}$. Then $\{\vec{r}_1\} \notin \mathcal{F}$, because \vec{s}_0 emulates \vec{r}_1 for \mathcal{F} and $f \downarrow_{\vec{r}_1/\vec{s}_0}$ maps the star $\{\vec{r}_1\} \subseteq \vec{S}_{\geq \vec{r}_1} \setminus \{\vec{r}_1\}$ to $\{\bar{s}_0\}$. Thus, \vec{r}_1 and \bar{r}_1 both lie outside $O^- \cup \vec{D}$, so $r_1 \in R$.

We can now hope to apply the induction hypothesis to $\mathcal{F}_1 := \mathcal{F} \cup \{\{\bar{r}_1\}\}$, because

$$O_1^- := O^- \cup \{\bar{r}_1\} = \{ \bar{s} \mid \{\bar{s}\} \in \mathcal{F}_1 \}$$

is again antisymmetric, and the set R_1 of separations in $S \setminus D$ with neither orientation in O_1^- is smaller than R . Also, \mathcal{F}_1 is a standard set of stars, because \mathcal{F} is and $r_1 \notin D$. But we still have to check that \vec{S} is \mathcal{F}_1 -separable.

To do so, consider (nontrivial and) nondegenerate separations $\vec{r}, \vec{r}' \in \vec{S}$ not forced by \mathcal{F}_1 such that $\vec{r} \leq \vec{r}'$. We have to find an $s_1 \in S$ with an orientation \vec{s}_1 emulates \vec{r} in \vec{S} for \mathcal{F}_1 and such that \bar{s}_1 emulates \vec{r}' in \vec{S} for \mathcal{F}_1 . By assumption, there is such an $s_1 \in S$ for \vec{r} and \vec{r}' with respect to \mathcal{F} ; let us take this s_1 , with orientations \vec{s}_1, \bar{s}_1 such that \vec{s}_1 emulates \vec{r} for \mathcal{F} and \bar{s}_1 emulates \vec{r}' for \mathcal{F} . We have to show that this emulation extends to \mathcal{F}_1 , i.e., that for the unique star $\{\bar{r}_1\}$ in $\mathcal{F}_1 \setminus \mathcal{F}$ we have $\{f \downarrow_{\vec{r}_1/\vec{s}_1}(\bar{r}_1)\} \in \mathcal{F}_1$ if $\vec{r} \leq \bar{r}_1 \neq \vec{r}$ (so that $\{\bar{r}_1\} \subseteq \vec{S}_{\geq \vec{r}} \setminus \{\vec{r}\}$), and $\{f \downarrow_{\vec{r}_1/\bar{s}_1}(\bar{r}_1)\} \in \mathcal{F}_1$ if $\vec{r}' \leq \bar{r}_1 \neq \vec{r}'$. In either case, the image \bar{s} of \bar{r}_1 under the relevant map is either equal to \bar{r}_1 (in which case we are done) or greater, by definition of the shift operator \downarrow . If $\bar{r}_1 < \bar{s}$, then $s \notin D$, since otherwise $\bar{r}_1 < \bar{s} = \bar{s}$ would be trivial and hence in O^- . And $\bar{s} < \bar{r}_1$ by (1), so $\bar{s} \in O^- \cup \vec{D}$ by the minimality of \bar{r}_1 in $\vec{S} \setminus (O^- \cup \vec{D})$. Thus $\bar{s} \in O^-$, and hence $\{\bar{s}\} \in \mathcal{F} \subseteq \mathcal{F}_1$, by the definition of O^- . This completes our proof that \vec{S} is \mathcal{F}_1 -separable.

We can thus apply the induction hypothesis to \mathcal{F}_1 . If it returns an \mathcal{F}_1 -tangle of S , then this is our desired \mathcal{F} -tangle. So we may assume that it returns an S -tree (T_1, α_1) over \mathcal{F}_1 . If this S -tree is even over \mathcal{F} , our proof is complete. We may thus assume that T_1 has a leaf x_1 associated with $\{\bar{r}_1\}$. We now apply Lemma 2.4 to prune and contract (T_1, α_1) to an irredundant S -tree over \mathcal{F}_1 that still contains x_1 , and in which no oriented edge other than the edge \vec{e} issuing from x_1 maps to \bar{r}_1 . For simplicity, let us continue to call this S -tree (T_1, α_1) . Let $\alpha'_1 := f \downarrow_{\vec{r}_1/\vec{s}_0} \circ \alpha_1$. By Lemma 4.2, (T_1, α_1) shifts to an S -tree (T_1, α'_1) over $\mathcal{F} \cup \{\{\bar{s}_0\}\}$, in which the star $\{\bar{s}_0\}$ is associated with x_1 but with no other leaf of T_1 . All the other nodes of T_1 are therefore associated with stars in \mathcal{F} .

If $\{\bar{s}_0\} \in \mathcal{F}$, then (T_1, α'_1) is in fact an S -tree over \mathcal{F} , completing our proof.

¹⁵It can happen that $\vec{r}_1 = \vec{r}_2$, in which case $r_1 = s_0 = r_2$.

We may thus assume that $\{\bar{s}_0\} \notin \mathcal{F}$, or equivalently that $\bar{s}_0 \notin O^-$. We can now use the induction hypothesis exactly as above (where we assumed that $\bar{s}_0 \notin O^-$), considering \bar{r}_2 in the same way as we just treated \bar{r}_1 , to obtain an irredundant S -tree (T_2, α'_2) over $\mathcal{F} \cup \{\{\bar{s}_0\}\}$ in which $\{\bar{s}_0\}$ is associated with a unique leaf x_2 , and all the other nodes are associated with stars in \mathcal{F} .

These trees can now be combined to the desired S -tree (T, α) over \mathcal{F} as in the proof of Theorem 3.1: add to the disjoint union $(T_2 - x_2) \cup (T_1 - x_1)$ the edge $y_2 y_1$ between the neighbour y_2 of x_2 in T_2 and the neighbour y_1 of x_1 in T_1 , put $\alpha(y_2, y_1) := \bar{s}_0$ and $\alpha(y_1, y_2) := \bar{s}_0$, and otherwise let α extend α'_1 and α'_2 . \square

5 Essential S -trees and \mathcal{F} -tangles: a refinement

Let us return to the question of how much of a restriction is our condition in the premise of the strong duality theorem that \mathcal{F} must be standard for \vec{S} , i.e., contain all co-trivial singletons, the stars $\{\bar{s}\}$ for which \bar{s} is trivial in \vec{S} . As noted before, any consistent orientation of S , and hence any \mathcal{F} -tangle, will contain all trivial separations and hence avoid all these singletons. So adding them to \mathcal{F} will not change the set of \mathcal{F} -tangles.

But neither would removing them. Which thus seems like a good idea, if only to avoid unnecessary clutter.

Removing the co-trivial singletons from \mathcal{F} would, however, have an impact on the set of S -trees over \mathcal{F} . The leaves of an S -tree over a standard \mathcal{F} can be associated with any co-trivial singleton star, but if we remove these stars from \mathcal{F} then such an S -tree will no longer be over \mathcal{F} .

We might try to repair this by removing those leaves from our S -tree (T, α) . The edge which such a leaf sends to its neighbour t , however, maps to a separation \bar{r} that would then be missing from the star $\alpha(\vec{F}_t) \in \mathcal{F}$ associated with t , perhaps knocking it out of \mathcal{F} . But as \bar{r} is trivial, its membership in $\alpha(\vec{F}_t)$ will not be the reason why we put $\alpha(\vec{F}_t)$ in \mathcal{F} in the first place: if the role of an S -tree over \mathcal{F} is to witness the nonexistence of an \mathcal{F} -tangle, then only the nontrivial separations in its stars are essential for that role. So let's try to delete all trivial separations from stars in \mathcal{F} , and see if we can retain an S -tree over the modified \mathcal{F} .

Given a separation system $(\vec{S}, \leq, *)$ and $\mathcal{F} \subseteq 2^{\vec{S}}$, define the *essential core* $\vec{\mathcal{F}}$ of \mathcal{F} as

$$\vec{\mathcal{F}} := \{F \setminus \vec{S}^- \mid F \in \mathcal{F}\},$$

where $\vec{S}^- \subseteq \vec{S}$ is the set of all separations that are trivial in \vec{S} . Note that if \mathcal{F} is standard then so is $\vec{\mathcal{F}}$, since inverses of trivial separations are never trivial. Let us call an S -tree (T, α) *essential* if it is irredundant, $\alpha(\vec{E}(T))$ contains no trivial separation, and all the sets $\alpha(\vec{F}_t)$ are antisymmetric.

Theorem 5.1. [3] *Let $(\vec{S}, \leq, *)$ a separation system, and $\mathcal{F} \subseteq 2^{\vec{S}}$ a set of stars.*

- (i) *The $\vec{\mathcal{F}}$ -tangles of S are precisely its \mathcal{F} -tangles.*
- (ii) *If (T, α) is any S -tree over \mathcal{F} , there is an essential S -tree (T', α') over $\vec{\mathcal{F}}$ such that T' is a minor of T and $\alpha' = \alpha \upharpoonright \vec{E}(T')$. Conversely, from any essential S -tree over $\vec{\mathcal{F}}$ we can obtain an S -tree over \mathcal{F} by adding leaves, if \mathcal{F} is standard for \vec{S} .*

Proof. (i) is immediate from the fact that \mathcal{F} -tangles, being consistent, contain all trivial separations and hence also avoid $\overline{\mathcal{F}}$.

For the first statement in (ii), let us start by making the given S -tree (T, α) irredundant by pruning it, as in Lemma 2.1. We then contract edges violating the desired antisymmetry of stars $\alpha(\overline{F}_t)$, as explained before Lemma 2.4. We finally make the resulting tree essential by deleting all its edges that α maps to trivial separations. This can be done recursively by deleting leaves associated with a co-trivial singleton: since \mathcal{F} consists of stars, Lemma 2.2 implies that any S -tree over \mathcal{F} with an edge mapping to a trivial separation will also have such an edge issuing from a leaf. (Recall that if \overline{s} is trivial then so is every $\overline{r} \leq \overline{s}$.) Pruning off leaves recursively in this way will leave a well-defined tree at the end, which has the properties desired for (T', α') .

For the second statement in (ii), let (T, α) be an essential S -tree over $\overline{\mathcal{F}}$, and consider a node $t \in T$. As $\alpha(\overline{F}_t) \in \overline{\mathcal{F}}$, there exists $F \in \mathcal{F}$ such that $F \setminus \alpha(\overline{F}_t)$ is a set of trivial separations \overline{s} . For each of these add a new leaf, joining it to t by an edge \overline{e} with $\alpha(\overline{e}) := \overline{s}$. \square

Theorem 5.1 allows us to strengthen each of the two alternatives in the strong duality theorem from its current version with the given set \mathcal{F} of stars to an ‘essential’ version with $\overline{\mathcal{F}}$ instead. So why didn’t we prove this stronger version directly?

The answer is pragmatic: this would have been possible, and we shall indicate in a moment how to do it. But it would have made the proof notationally more technical. As the proof stands, we need to allow inessential S -trees, because they can arise in the induction step when we combine two shifted S -trees, even if these were essential before the shift.

Indeed, recall what happens to the leaves of an S -tree when we shift it, by $f \downarrow_{\overline{s}_0}^{\overline{r}}$ say. A leaf, associated with $\{\overline{r}'\}$, say, where $r \neq r' \neq s_0$ for simplicity, will be associated in the shifted tree with the star $\{\overline{r}' \vee \overline{s}_0\}$, because $\overline{r} \leq \overline{r}'$. But if $\overline{r}' < \overline{s}_0$, as will frequently happen, then this star is a co-trivial singleton, because $\overline{r}' \wedge \overline{s}_0 < \overline{r}'$ as well as $\overline{r}' \wedge \overline{s}_0 < \overline{s}_0 < \overline{r}'$.

The way to overcome this problem is indicated in Footnote 14, with S' the set of nontrivial separations in S . When we shift a star $\sigma \in \mathcal{F}$, its shift $\alpha(\sigma)$ may contain trivial separations, but we could simply delete these to make the shifted S -tree essential, as in the proof of Theorem 5.1(ii). To ensure that it is again over \mathcal{F} , we would need to replace the current requirement in the definition of \mathcal{F} -separable, that $f \downarrow_{\overline{s}_0}^{\overline{r}}$ should map to \mathcal{F} any $\sigma \in \mathcal{F}$ that contains a separation $\overline{s} \geq \overline{r}$, with the requirement that for any such $\sigma \in \mathcal{F}$ we have $f \downarrow_{\overline{s}_0}^{\overline{r}}(\sigma) \cap \overline{S}' \in \mathcal{F}$ if $f \downarrow_{\overline{s}_0}^{\overline{r}}(\overline{s}) \in \overline{S}'$.

Acknowledgement

We enjoyed a visit by Frédéric Mazoit to the University of Hamburg in the summer of 2013, when he gave a series of lectures on the bramble-based duality theorems for width parameters developed in [1, 9]. We then tried, unsuccessfully, to apply these to obtain a duality theorem for k -blocks. This triggered the development of a similar unified duality theory for tangles instead of brambles, first in graphs and later more generally. Our original goal, a duality theorem for k -blocks, was eventually achieved in [5], as an application of Theorem 4.3.

References

- [1] O. Amini, F. Mazoit, N. Nisse, and S. Thomassé. Submodular partition functions. *Discrete Appl. Math.*, 309(20):6000–6008, 2009.
- [2] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. k -Blocks: a connectivity invariant for graphs. *SIAM J. Discrete Math.*, 28:1876–1891, 2014.
- [3] R. Diestel. Tree sets. arXiv:1512.03781, 2015.
- [4] R. Diestel. *Graph Theory* (5th edition). Springer-Verlag, 2017. Electronic edition available at <http://diestel-graph-theory.com/>.
- [5] R. Diestel, J. Erde, and P. Eberenz. Duality theorem for tangles in graphs. arXiv:1605.09139, 2016.
- [6] R. Diestel and S. Oum. Tangle-tree duality in graphs and matroids. ArXiv, 2017.
- [7] R. Diestel and G. Whittle. Tangles and the Mona Lisa. arXiv:1603.06652, 2016.
- [8] J. Geelen, B. Gerards, N. Robertson, and G. Whittle. Obstructions to branch-decomposition of matroids. *J. Combin. Theory (Series B)*, 96:560–570, 2006.
- [9] L. Lyaudet, F. Mazoit, and S. Thomassé. Partitions versus sets: a case of duality. *European J. Combin.*, 31(3):681–687, 2010.
- [10] W. Mader. Über n -fach zusammenhängende Eckenmengen in Graphen. *J. Combin. Theory (Series B)*, 25:74–93, 1978.
- [11] B. Reed. Tree width and tangles: a new connectivity measure and some applications. In R. Bailey, editor, *Surveys in Combinatorics*. Cambridge Univ. Press, 1997.
- [12] N. Robertson and P. Seymour. Graph minors. V. Excluding a planar graph. *J. Combin. Theory (Series B)*, 41:92–114, 1986.
- [13] N. Robertson and P. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)*, 52:153–190, 1991.
- [14] N. Robertson and P. Seymour. Graph minors. XVI. Excluding a non-planar graph. *J. Combin. Theory (Series B)*, 89:43–76, 2003.
- [15] P. Seymour and R. Thomas. Call routing and the ratcatcher. *Combinatorica*, 14(2):217–241, 1994.