

Duality theorems for blocks and tangles in graphs

Reinhard Diestel, Philipp Eberenz and Joshua Erde

July 6, 2017

Abstract

We prove a duality theorem applicable to a wide range of specialisations, as well as to some generalisations, of tangles in graphs. It generalises the classical tangle duality theorem of Robertson and Seymour, which says that every graph either has a large-order tangle or a certain low-width tree-decomposition witnessing that it cannot have such a tangle.

Our result also yields duality theorems for profiles and for k -blocks. This solves a problem studied, but not solved, by Diestel and Oum and answers an earlier question of Carmesin, Diestel, Hamann and Hundertmark.

1 Introduction

There are a number of theorems about the structure of sparse graphs that assert a duality between the existence of a highly connected substructure and a tree-like overall structure. For example, if a connected graph G has no 2-connected subgraph or minor, it is a tree. Less trivially, the graph minor structure theorem of Robertson and Seymour says that if G has no K_n -minor then it has a tree-decomposition into parts that are ‘almost’ embeddable in a surface of bounded genus; see [5].

Another example of a highly connected substructure is that of a k -block, introduced by Mader [12] in 1978 and studied more recently in [1, 2, 3]. This is a maximal set of at least k vertices in a graph G such that no two of them can be separated in G by fewer than k vertices.

One of our main results is that the non-existence of a k -block, too, is always witnessed by a tree structure (Theorem 2). This problem was raised in [1, Sec. 7]. In [9], Diestel and Oum used a new theory of ‘abstract separation systems’ [4] in pursuit of this problem, but were unable to find the tree structures needed: the simplest witnesses to the nonexistence of k -blocks they could find are described in [8], but they are more complicated than trees. Our proof of Theorem 2, and the rest of this paper, are still based on the theory developed in [4] and [9], and we show that there are tree-like obstructions to the existence of k -blocks after all.

Tangles, introduced by Robertson and Seymour in [13], are substructures of graphs that also signify high local connectivity, but of a less tangle-like kind than subgraphs, minors, or blocks. Basically, a tangle does not tell us ‘what’ that substructure is, but only ‘where’ it is: by orienting all the low-order separations of the graph in some consistent way, which we then think of as pointing ‘towards the tangle’. See Section 2.2 below, or [5], for a formal introduction to tangles.

Tangles come with dual tree structures called *branch decompositions*. Although defined differently, they can be thought of as tree-decompositions of a

particular kind. In [9, 10], Diestel and Oum generalised the notion of tangles to ways of consistently orienting the low-order separations of a graph so as to describe other known types of highly connected substructures too, such as those dual to tree-decompositions of low width. We shall build on [9] to find dual tree structures for various types of tangles, for blocks, and for ‘profiles’: a common generalisation of blocks and tangles introduced in [7] and defined formally in Section 2.3.

Let us describe our results more precisely. A classical k -tangle, as in [13], is an orientation of all the separations $\{A, B\}$ of order $< k$ in a graph G , say as (A, B) rather than as (B, A) , so that no three of these cover G by the subgraphs that G induces on their ‘small sides’ A . Let \mathcal{T} denote the set of all such forbidden triples of oriented separations of G , irrespective of their order. The tangle duality theorem of Robertson and Seymour then asserts that, given k , either the set S_k of all the separations of G of order $< k$ can be oriented in such a way as to induce no triple from \mathcal{T} – an orientation of S_k we shall call a \mathcal{T} -tangle – or G has a tree-decomposition of a particular type: one from which it is clear that G cannot have a k -tangle, i.e., a \mathcal{T} -tangle of S_k .

Now consider any superset \mathcal{F} of \mathcal{T} . Given k , the orientations of S_k with no subset in \mathcal{F} will then be particular types of tangles. Our \mathcal{F} -tangle duality theorem yields duality theorems for all these: if G contains no such ‘special’ tangle, it will have a tree-decomposition that witnesses this. Formally, to every such \mathcal{F} and k there will correspond a class $\mathcal{T}_{\mathcal{F}}(k)$ of tree-decompositions that witness the non-existence of an \mathcal{F} -tangle of S_k , and which are shown to exist whenever a graph has no \mathcal{F} -tangle of S_k :

Theorem 1. *For every finite graph G , every set $\mathcal{F} \supseteq \mathcal{T}$ of sets of separations of G , and every integer $k > 2$, exactly one of the following statements holds:*

- G admits an \mathcal{F} -tangle of S_k ;
- G has a tree-decomposition in $\mathcal{T}_{\mathcal{F}}(k)$.

Every k -block also defines an orientation of S_k : as no separation $\{A, B\} \in S_k$ separates it, it lies entirely in A or entirely in B . These orientations of S_k need not be k -tangles, so we cannot apply Theorem 1 to obtain a duality theorem for k -blocks. But still, we shall be able to define classes $\mathcal{T}_{\mathcal{B}}(k)$ of tree-decompositions that witness the non-existence of a k -block, in the sense that graphs with such a tree-decomposition cannot contain one, and which always exist for graphs without a k -block:

Theorem 2. *For every finite graph G and every integer $k > 0$ exactly one of the following statements holds:*

- G contains a k -block;
- G has a tree-decomposition in $\mathcal{T}_{\mathcal{B}}(k)$.

Finally, we define classes $\mathcal{T}_{\mathcal{P}}(k)$ of tree-decompositions which graphs with a k -profile cannot have, and prove the following duality theorem for profiles:

Theorem 3. *For every finite graph G and every integer $k > 2$ exactly one of the following statements holds:*

- G has a k -profile;
- G has a tree-decomposition in $\mathcal{T}_{\mathcal{P}}(k)$.

For readers already familiar with profiles [7] we remark that, in fact, we shall obtain a more general result than Theorem 3: our Theorem 15 is a duality theorem for all regular profiles in arbitrary submodular abstract separation systems, including the standard ones in graphs and matroids but many others too [11].

Like Theorem 1, Theorems 2 and 3 are ‘structural’ duality theorems in that they identify a structure that a graph G cannot have if it contains a k -block or k -profile, and must have if it does not. Alternatively, we can express the same duality more compactly in terms of graph invariants, as follows. Let

$$\begin{aligned}\beta(G) &:= \max \{ k \mid G \text{ has a } k\text{-block} \} \\ \pi(G) &:= \max \{ k \mid G \text{ has a } k\text{-profile} \}\end{aligned}$$

be the *block number* and the *profile number* of G , respectively, and let

$$\begin{aligned}\text{bw}(G) &:= \min \{ k \mid G \text{ has a tree-decomposition in } \mathcal{T}_{\mathcal{B}}(k+1) \} \\ \text{pw}(G) &:= \min \{ k \mid G \text{ has a tree-decomposition in } \mathcal{T}_{\mathcal{P}}(k+1) \}\end{aligned}$$

be its *block-width* and *profile-width*. Theorems 2 and 3 can now be rephrased as

Corollary 4. *As invariants of finite graphs, the block and profile numbers agree with the block- and profile-widths:*

$$\beta = \text{bw} \quad \text{and} \quad \pi = \text{pw}. \quad \square$$

In Section 2 we introduce just enough about abstract separation systems [4] to state the fundamental duality theorem of [9], on which all our proofs will be based. In Section 3 we give a proof of our main result, a duality theorem for regular profiles in submodular abstract separation systems. In Section 4 we apply this to obtain structural duality theorems for k -blocks and k -profiles, and deduce Theorems 1–3 as corollaries. In Section 5 we derive some bounds for the above width-parameters in terms of tree-width and branch-width.

Any terms or notation left undefined in this paper are explained in [5].

2 Background Material

2.1 Separation systems

A *separation* of a graph G is a set $\{A, B\}$ of subsets of $V(G)$ such that $A \cup B = V$, and there is no edge of G between $A \setminus B$ and $B \setminus A$. There are two *oriented separations* associated with a separation, (A, B) and (B, A) . Informally we think of (A, B) as *pointing towards* B and *away from* A . We can define a partial ordering on the set of oriented separations of G by

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

The *inverse* of an oriented separation (A, B) is the separation (B, A) , and we note that mapping every oriented separation to its inverse is an involution which reverses the partial ordering.

In [9] Diestel and Oum generalised these properties of separations of graphs and worked in a more abstract setting. They defined a *separation system* $(\vec{S}, \leq, *)$ to be a partially ordered set \vec{S} with an order-reversing involution $*$.

The elements of \vec{S} are called *oriented separations*. Often a given element of \vec{S} is denoted by \vec{s} , in which case its inverse \vec{s}^* will be denoted by \overleftarrow{s} , and vice versa. Since $*$ is ordering reversing we have that, for all $\vec{r}, \vec{s} \in S$,

$$\vec{r} \leq \vec{s} \text{ if and only if } \overleftarrow{r} \geq \overleftarrow{s}.$$

A *separation* is a set of the form $\{\vec{s}, \overleftarrow{s}\}$, and will be denoted by simply s . The two elements \vec{s} and \overleftarrow{s} are the *orientations* of s . The set of all such pairs $\{\vec{s}, \overleftarrow{s}\} \subseteq \vec{S}$ will be denoted by S . If $\vec{s} = \overleftarrow{s}$ we say s is *degenerate*. Conversely, given a set $S' \subseteq S$ of separations we write $\vec{S}' := \bigcup S'$ for the set of all orientations of its elements. With the ordering and involution induced from \vec{S} , this will form a separation system.

Given a separation of a graph $\{A, B\}$ we can identify it with the pair $\{(A, B), (B, A)\}$ and in this way any set of oriented separations in a graph which is closed under taking inverses forms a separation system. When we refer to an oriented separation in a context where the notation explicitly indicates orientation, such as \vec{s} or (A, B) , we will usually suppress the prefix ‘‘oriented’’ to improve the flow of the narrative.

The *separator* of a separation $s = \{A, B\}$ in a graph, and of its orientations \vec{s} , is the set $A \cap B$. The *order* of s and \vec{s} , denoted as $|s|$ or as $|\vec{s}|$, is the cardinality of the separator, $|A \cap B|$. Note that if $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$ are separations then so are their *corner separations* $\vec{r} \vee \vec{s} := (A \cup C, B \cap D)$ and $\vec{r} \wedge \vec{s} := (A \cap C, B \cup D)$. Our function $\vec{s} \mapsto |\vec{s}|$ is clearly *symmetric* in that $|\vec{s}| = |\overleftarrow{s}|$, and *submodular* in that

$$|\vec{r} \vee \vec{s}| + |\vec{r} \wedge \vec{s}| \leq |\vec{r}| + |\vec{s}|$$

(in fact, with equality).




Figure 1: The corner separation $\vec{r} \vee \vec{s} = (A \cup C, B \cap D)$

If an abstract separation system $(\vec{S}, \leq, *)$ forms a lattice, i.e., if there exist binary operations \vee and \wedge on \vec{S} such that $\vec{r} \vee \vec{s}$ is the supremum and $\vec{r} \wedge \vec{s}$ is the infimum of \vec{r} and \vec{s} , then we call $(\vec{S}, \leq, *, \vee, \wedge)$ a *universe* of (oriented) separations. By (2.1), it satisfies De Morgan’s law:

$$(\vec{r} \vee \vec{s})^* = \overleftarrow{r} \wedge \overleftarrow{s}. \tag{2.1}$$

Any real, non-negative, symmetric and submodular function on a universe of separations, usually denoted as $\vec{s} \mapsto |\vec{s}|$, will be called an *order function*.

Two separations r and s are *nested* if they have \leq -comparable orientations. Two oriented separations \vec{r} and \vec{s} are *nested* if r and s are nested.¹ We say that \vec{r} *points towards* s (and \overleftarrow{r} *points away from* s) if $\vec{r} \leq \vec{s}$ or $\vec{r} \leq \overleftarrow{s}$. So two nested oriented separations are either \leq -comparable, or they point towards each other, or they point away from each other. If \vec{r} and \vec{s} are not nested we say that they *cross*. A set of separations S is *nested* if every pair of separations in S is nested, and a separation s is *nested* with a nested set of separations S if $S \cup \{s\}$ is nested.

A separation $\vec{r} \in \vec{S}$ is *trivial in* \vec{S} , and \overleftarrow{r} is *co-trivial*, if there exist an $s \in S$ such that $\vec{r} < \vec{s}$ and $\vec{r} < \overleftarrow{s}$. Note that if \vec{r} is trivial, witnessed by some s , then, since the involution on \vec{S} is order-reversing, we have $\vec{r} < \vec{s} < \overleftarrow{r}$. So, in particular, \overleftarrow{r} cannot also be trivial. Separations \vec{s} such that $\vec{s} \leq \overleftarrow{s}$, trivial or not, will be called *small*.

In the case of separations of a graph (V, E) , the small separations are precisely those of the form (A, V) . The trivial separations are those of the form (A, V) with $A \subseteq C \cap D$ for some separation $\{C, D\} \neq \{A, B\}$. Finally we note that there is only one degenerate separation in a graph, (V, V) .

2.2 Tangle-tree duality in separation systems

Let \vec{S} be a separation system. An *orientation* of S is a subset $O \subseteq \vec{S}$ which for each $s \in S$ contains exactly one of its orientations \vec{s} or \overleftarrow{s} . Given a universe \vec{U} of separations with an order function, such as all the oriented separations of a given graph, we denote by

$$\vec{S}_k = \{\vec{s} \in \vec{U} : |\vec{s}| < k\}$$

the set of all its separations of order less than k . Note that \vec{S}_k is again a separation system. But it is not necessarily a universe, since it may fail to be closed under the operations \vee and \wedge .

If we have some structure \mathcal{C} in a graph that is ‘highly connected’ in some sense, we should expect that no low order separation will divide it: that is, for every separation s of sufficiently low order, \mathcal{C} should lie on one side of s but not the other. Then \mathcal{C} will *orient* s as \vec{s} or \overleftarrow{s} , choosing the orientation that ‘points to where it lies’ according to some convention. For graphs, our convention is that the orientated separation (A, B) points towards B . And that if \mathcal{C} is a K_n -minor of G with $n \geq k$, say, then \mathcal{C} ‘lies on the side B ’ if it has a branch set in $B \setminus A$. (Note that it cannot have a branch set in $A \setminus B$ then.) Then \mathcal{C} orients $\{A, B\}$ towards B by choosing (A, B) rather than (B, A) . In this way, \mathcal{C} induces an orientation of all of S_k .

The idea of [9], now, following the idea of tangles, was to *define* ‘highly connected substructures’ in this way: as orientations of a given set S of separations.

Any concrete example of ‘highly-connected substructures’ in a graph, such as a K_n -minor or a k -block, will not induce arbitrary orientations of S_k : these orientations will satisfy some consistency rules. For example, consider two separations $(A, B) < (C, D)$. If our ‘highly connected’ structure \mathcal{C} orients $\{C, D\}$ towards D then, since $B \supseteq D$ it should not orient $\{A, B\}$ towards A .

¹In general we will use terms defined for separations informally for oriented separations when the meaning is clear, and vice versa.

We call an orientation O of a set S of separations in some universe \vec{U} *consistent* if whenever we have distinct r and s such that $\vec{r} < \vec{s}$, the set O does not contain both \overleftarrow{r} and \overleftarrow{s} . Note that a consistent orientation of S must contain all separations \vec{r} that are trivial in S since, if $\vec{r} < \vec{s}$ and $\vec{r} < \overleftarrow{s}$, then \overleftarrow{r} would be inconsistent with whichever orientation of s lies in O .

Given a set \mathcal{F} , we say that an orientation O of S *avoids* \mathcal{F} if there is no $F \in \mathcal{F}$ such that $F \subseteq O$. So for example an orientation of S is consistent if it avoids $\mathcal{F} = \{\{\overleftarrow{r}, \overleftarrow{s}\} \subseteq \vec{S} : r \neq s, \vec{r} < \vec{s}\}$. In general we will define the highly connected structures we consider by the collection \mathcal{F} of subsets they avoid. For example a *tangle of order k* , or *k -tangle*, in a graph G is an orientation of S_k which avoids the set of triples

$$\mathcal{T} = \{\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\} \subseteq \vec{U} : \bigcup_{i=1}^3 G[A_i] = G\}. \quad (2.2)$$

Here, the three separations need not be distinct, so any \mathcal{T} -avoiding orientation of S_k will be consistent. More generally, we say that a consistent orientation of a set S of separations which avoids some given set \mathcal{F} is an \mathcal{F} -*tangle* (of S).

Given a set S of separations, an S -*tree* is a pair (T, α) , of a tree T and a function $\alpha : \overrightarrow{E(T)} \rightarrow \vec{S}$ from the set $\overrightarrow{E(T)}$ of directed edges of T such that

- For each edge $(t_1, t_2) \in \overrightarrow{E(T)}$, if $\alpha(t_1, t_2) = \vec{s}$ then $\alpha(t_2, t_1) = \overleftarrow{s}$.

The S -tree is said to be *over* a set \mathcal{F} if

- For each vertex $t \in T$, the set $\{\alpha(t', t) : (t', t) \in \overrightarrow{E(T)}\}$ is in \mathcal{F} .

Particularly interesting classes of S -trees are those over sets \mathcal{F} of ‘stars’. A set σ of nondegenerate oriented separations is a *star* if $\vec{r} \leq \overleftarrow{s}$ for all distinct $\vec{r}, \overleftarrow{s} \in \sigma$. We say that a set \mathcal{F} *forces* a separation \vec{r} if $\{\overleftarrow{r}\} \in \mathcal{F}$. And \mathcal{F} is *standard* if it forces every trivial separation in \vec{S} .

The main result of [9] asserts a duality between S -trees over \mathcal{F} and \mathcal{F} -tangles when \mathcal{F} is a standard set of stars satisfying a certain closure condition. Let us describe this next.

Suppose we have a separation \vec{r} which is neither trivial nor degenerate. Let $S_{\geq \vec{r}}$ be the set of separations $x \in S$ that have an orientation $\vec{x} \geq \vec{r}$. Given $x \in S_{\geq \vec{r}} \setminus \{r\}$ we have, since \vec{r} is nontrivial, that only one of the two orientations of x , say \vec{x} , is such that $\vec{x} \geq \vec{r}$ and x is not degenerate. For every $\overleftarrow{s} \geq \vec{r}$ we can define a function $f \downarrow_{\overleftarrow{s}}^{\vec{r}}$ on $\vec{S}_{\geq \vec{r}} \setminus \{\overleftarrow{r}\}$ by²

$$f \downarrow_{\overleftarrow{s}}^{\vec{r}}(\vec{x}) := \vec{x} \vee \overleftarrow{s} \text{ and } f \downarrow_{\overleftarrow{s}}^{\vec{r}}(\overleftarrow{x}) := (\vec{x} \vee \overleftarrow{s})^*.$$

In general, the image in \vec{U} of this function need not lie in \vec{S} .

We say that $\overleftarrow{s} \in \vec{S}$ *emulates* $\vec{r} \in \vec{S}$ in \vec{S} if $\vec{r} \leq \overleftarrow{s}$ and the image of $f \downarrow_{\overleftarrow{s}}^{\vec{r}}$ is contained in \vec{S} . Given a standard set \mathcal{F} of stars, we say further that \overleftarrow{s} *emulates* \vec{r} in \vec{S} for \mathcal{F} if \overleftarrow{s} emulates \vec{r} in \vec{S} and the image under $f \downarrow_{\overleftarrow{s}}^{\vec{r}}$ of every star $\sigma \subseteq \vec{S}_{\geq \vec{r}} \setminus \{\overleftarrow{r}\}$ that contains some separation \vec{x} with $\vec{x} \geq \vec{r}$ is again in \mathcal{F} .

²The exclusion of \overleftarrow{r} here is for a technical reason: if $\vec{r} < \overleftarrow{r}$, we do not want to define $f \downarrow_{\overleftarrow{s}}^{\vec{r}}(\overleftarrow{r})$ explicitly, but implicitly as the inverse of $f \downarrow_{\overleftarrow{s}}^{\vec{r}}(\vec{r})$.




Figure 2: Shifting a separation $\vec{x} \geq \vec{r}$ to $f \downarrow_{\vec{s}}^{\vec{r}}(\vec{x}) = \vec{x} \vee \vec{s}$.

We say that a separation system \vec{S} is *separable* if for any two nontrivial and nondegenerate separations $\vec{r}, \vec{r}' \in \vec{S}$ such that $\vec{r} \leq \vec{r}'$ there exists a separation $s \in S$ such that \vec{s} emulates \vec{r} in \vec{S} and \overleftarrow{s} emulates \vec{r}' in \vec{S} . We say 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 such that $\vec{r} \leq \vec{r}'$ there exists a separation $s \in S$ with an orientation \vec{s} that emulates \vec{r} in \vec{S} for \mathcal{F} and such that \overleftarrow{s} emulates \vec{r}' in \vec{S} for \mathcal{F} . Often one proves that \vec{S} is \mathcal{F} -separable in two steps, by first showing that it is separable, and then showing that \mathcal{F} is *closed under shifting*: that whenever \vec{s} emulates (in \vec{S}) some nontrivial and nondegenerate \vec{r} not forced by \mathcal{F} , then it does so for \mathcal{F} .

We are now in a position to state the Strong Duality Theorem from [9].

Theorem 5. *Let \vec{S} be a separation system in some universe of separations, and \mathcal{F} a standard set of stars. If \vec{S} is \mathcal{F} -separable, exactly one of the following assertions holds:*

- *There exists an S -tree over \mathcal{F} ;*
- *There exists an \mathcal{F} -tangle of S .*

The property of being \mathcal{F} -separable may seem a rather strong condition. However in [10] it is shown that for every graph the set \vec{S}_k is separable, and all the sets \mathcal{F} of stars whose exclusion describes classical notions of highly connected substructures are closed under shifting. Hence in all these cases \vec{S}_k is \mathcal{F} -separable, and Theorem 5 applies.

One of our main tasks will be to extend the applicability of Theorem 5 to sets \mathcal{F} of separations that are not stars, by constructing a related set \mathcal{F}^* of stars whose exclusion is tantamount to excluding \mathcal{F} .

2.3 Blocks, tangles, and profiles

Suppose we have a graph $G = (V, E)$ and are considering the set U of its separations. As mentioned before, it is easy to see that the tangles of order k in G , as defined by Robertson and Seymour [13], are precisely the \mathcal{T} -tangles of $S_k = \{s \in U : |s| < k\}$. In this case, if we just consider the set of stars in \mathcal{T} ,

$$\begin{aligned} \mathcal{T}^* := \{ \{ (A_1, B_1), (A_2, B_2), (A_3, B_3) \} \subseteq \vec{S} : \\ (A_i, B_i) \leq (B_j, A_j) \text{ for all } i, j \text{ and } \bigcup_{i=1}^3 G[A_i] = G \}, \end{aligned}$$

then the \mathcal{T}^* -tangles of S_k are precisely its \mathcal{T} -tangles [10]. That is, a consistent orientation of S_k avoids \mathcal{T} if and only if it avoids \mathcal{T}^* . It is a simple check that \mathcal{T}^* is a standard set of stars which is closed under shifting, and hence Theorem 5 tells us that every graph either has a tangle of order k or an S_k -tree over \mathcal{T}^* , but not both.

Another highly connected substructure that has been considered recently in the literature are k -blocks. Given $k \in \mathbb{N}$ we say a set I of at least k vertices in a graph G is $(<k)$ -inseparable if no set Z of fewer than k vertices separates any two vertices of $I \setminus Z$ in G . A maximal $(<k)$ -inseparable set of vertices is called a k -block. These objects were first considered by Mader [12], but have been the subject of recent research [1, 2, 3].

As indicated earlier, every k -block b of G defines an orientation $O(b)$ of S_k :

$$O(b) := \{(A, B) \in \overrightarrow{S}_k : b \subseteq B\}.$$

Indeed, for each separation $\{A, B\} \in S_k$ exactly one of (A, B) and (B, A) will be in $O(b)$, since $A \cap B$ is too small to contain b and does not separate any two of its vertices. Hence, $O(b)$ is indeed an orientation of S_k . Note also that $O(b) \neq O(b')$ for distinct k -blocks $b \neq b'$: by their maximality as k -indivisible sets of vertices there exists a separation $\{A, B\} \in S_k$ such that $A \setminus B$ contains a vertex of b and $B \setminus A$ contains a vertex of b' , which implies that $(A, B) \in O(b')$ and $(B, A) \in O(b)$.

The orientations $O(b)$ of S_k defined by a k -block b clearly avoid

$$\mathcal{B}_k := \left\{ \{(A_i, B_i) : i \in I\} \subseteq \overrightarrow{U} : \left| \bigcap_{i \in I} B_i \right| < k \right\},$$

since $b \subseteq B_i$ for every $(A_i, B_i) \in O(b)$ and $|b| \geq k$. Also, it is easily seen that every $O(b)$ is consistent. Thus, every such orientation $O(b)$ is an \mathcal{F} -tangle of S_k for $\mathcal{F} = \mathcal{B}_k$. Conversely, if $O \subseteq S_k$ is a \mathcal{B}_k -tangle of S_k , then $b := \bigcap \{B \mid (A, B) \in O\}$ is easily seen to be a k -block, and $O = O(b)$. The orientations of S_k that are defined by a k -block, therefore, are precisely its \mathcal{B}_k -tangles.

The \mathcal{B}_k -tangles of S_k and its \mathcal{T} -tangles (i.e., the ordinary k -tangles of G) share the property that if they contain separations (A, B) and (C, D) , then they cannot contain the separation $(B \cap D, A \cup C)$. Indeed, clearly this condition is satisfied by $O(b)$ for any k -block b , since if $b \subseteq B$ and $b \subseteq D$ then $b \subseteq B \cap D$ and hence $b \not\subseteq A \cup C$ if $\{B \cap D, A \cup C\} \in S_k$. For tangles, suppose that some tangle contains such a triple $\{(A, B), (C, D), (B \cap D, A \cup C)\}$. Since $\{A, B\}$ and $\{C, D\}$ are separations of G , every edge not contained in $G[A]$ or $G[C]$ must be in $G[B]$ and $G[D]$, and hence in $G[B \cap D]$. Therefore $G[A] \cup G[C] \cup G[B \cap D] = G$, contradicting the fact that the tangle avoids \mathcal{T} .

Informally, if we think of the side of an oriented separation to which it points as ‘large’, then the orientations of S_k that form a tangle or are induced by a k -block have the natural property that if B is the large side of $\{A, B\}$ and D is the large side of $\{C, D\}$ then $B \cap D$ should be the large side of $\{A \cup C, B \cap D\}$ – if this separation is also in S_k , and therefore oriented by O . That is, the largeness of separation sides containing blocks or tangles is preserved by taking intersections.

Consistent orientations with this property are known as ‘profiles’. Formally, a k -profile in G is a \mathcal{P} -tangle of S_k where

$$\mathcal{P} := \left\{ \sigma \subseteq \overrightarrow{U} \mid \exists A, B, C, D \subseteq V : \sigma = \{(A, B), (C, D), (B \cap D, A \cup C)\} \right\}.$$

As we have seen,

Lemma 6. *All orientations of S_k that are tangles, or of the form $O(b)$ for some k -block b in G , are k -profiles in G . \square*

We remark that, unlike in the case of \mathcal{T} , the subset \mathcal{P}' consisting of just the stars in \mathcal{P} yields a wider class of tangles: there are \mathcal{P}' -tangles of S_k that are not \mathcal{P} -tangles, i.e., which are not k -profiles.

More generally, if \vec{S} is any separation system contained in some universe \vec{U} , we can define a *profile of S* to be any \mathcal{P} -tangle of S where

$$\mathcal{P} := \{ \sigma \subseteq \vec{U} \mid \exists \vec{r}, \vec{s} \in \vec{U} : \sigma = \{ \vec{r}, \vec{s}, \overleftarrow{r} \wedge \overleftarrow{s} \} \}.$$

In particular, all \mathcal{F} -tangles with $\mathcal{F} \supseteq \mathcal{P}$ will be profiles.

The initial aim of Diestel and Oum in developing their duality theory [9] had been to find a duality theorem broad enough to imply duality theorems for k -blocks and k -profiles. Although their theory gave rise to a number of unexpected results [10], a duality theorem for blocks and profiles was not among these; see [8] for a summary of their findings on this problem.

Our next goal is to show that their Strong Duality Theorem does implies duality theorems for blocks and profiles after all.

3 A duality theorem for abstract profiles

In this section we will show that Theorem 5 can be applied to many more types of profiles than originally thought. These will include both k -profiles and k -blocks in graphs.

We say that a separation system in some universe³ is *submodular* if for every two of its elements \vec{r}, \vec{s} it also contains at least one of $\vec{r} \wedge \vec{s}$ and $\vec{r} \vee \vec{s}$. Given any (submodular) order function on a universe, the separation system

$$\vec{S}_k = \{ \vec{r} \in \vec{U} : |\vec{r}| < k \}$$

is submodular for each k . In particular, for any graph G , its universe \vec{U} of separations and, for any integer $k \geq 1$, the separation system \vec{S}_k , is submodular.

We say that a subset O of U is *strongly consistent* if it does not contain both \overleftarrow{r} and \overleftarrow{s} for any $\vec{r}, \vec{s} \in \vec{S}$ with $\vec{r} < \vec{s}$ (but not necessarily $r \neq s$, as in the definition of ‘consistent’). An orientation O of S , therefore, is strongly consistent if and only if for every $\vec{s} \in O$ it also contains every $\vec{r} \leq \vec{s}$ with $r \in S$. In particular, then, O cannot contain any \vec{s} such that $\overleftarrow{s} \leq \vec{s}$ (i.e., with \overleftarrow{s} is small).

Let us call an orientation O of S *regular* if it contains all the small separations in \vec{S} .

Lemma 7. *An orientation O of a separation system \vec{S} is strongly consistent if and only if it is consistent and regular.*

³Although submodular separation systems \vec{S} have to lie in some universe \vec{U} in order for \wedge and \vee to be defined on \vec{S} (but with images that may lie in $\vec{U} \setminus \vec{S}$), the choice of \vec{U} , given \vec{S} , will not matter to us. We shall therefore usually introduce submodular separation systems \vec{S} without formally introducing such a universe $\vec{U} \supseteq \vec{S}$.

Proof. Clearly every strongly consistent orientation O is also consistent. Suppose some small $\vec{s} \in \vec{S}$ is not in O . Then $\overleftarrow{s} \in O$, since O is an orientation of S . Thus, $\vec{s} < \overleftarrow{s} \in O$. But this implies $\vec{s} \in O$, since O is strongly consistent, contradicting the choice of \vec{s} . Hence O contains every small separation.

Conversely, suppose O is a consistent orientation of S that is not strongly consistent. Then O contains two distinct oriented separations \overleftarrow{r} and \vec{s} such that $\overleftarrow{r} < \vec{s}$ and $r = s$. Thus, $\overleftarrow{s} = \overleftarrow{r} < \vec{s}$ is small but not in O , as $\vec{s} \in O$. Hence O does not contain all small separations in \vec{S} . \square

For example, the \mathcal{B}_k -tangles of S_k in a graph, as well as its ordinary tangles of order k , are regular by Lemma 7: they clearly contain all small separations in \vec{S}_k , those of the form (A, V) with $|A| < k$, since they cannot contain their inverses (V, A) . More generally:

Lemma 8. *For $k > 2$, every k -profile in a graph $G = (V, E)$ is regular.*

Proof. We have to show that every k -profile O in G contains every small separation in \vec{S}_k . Recall that these are precisely the separations (A, V) of G such that $|A| < k$.

Suppose first that $|A| < k - 1$. Let A' be any set such that $|A'| = k - 1$ and $A \subset A'$. Then $\{A', V\} \in S_k$, and $(A, V) < (A', V)$ as well as $(A, V) < (V, A')$. Since O contains (A', V) or (V, A') , its consistency implies that it also contains (A, V) .

If $|A| = k - 1$ then, since $k > 2$, we can pick two non-empty sets $A', A'' \subsetneq A$ such that $A' \cup A'' = A$. Since $|A'|, |A''| < k - 1$, by the preceding discussion both (A', V) and (A'', V) lie in O . As $(V, A) \wedge (V, A'') = (V \cap V, A' \cup A'') = (V, A)$ and $\{(A', V), (A'', V), (V, A)\} \in \mathcal{P}$, the fact that O is a profile implies that $(V, A) \notin O$, so again $(A, V) \in O$ as desired. \square

There can be exactly one irregular 1-profile in a graph $G = (V, E)$, and only if G is connected: the set $\{(V, \emptyset)\}$.

Graphs can also have irregular 2-profiles, but they are easy to characterise. Indeed, consider a 2-profile O and small separation $(\{x\}, V)$. Suppose first that x is a cutvertex of G , in the sense that there exists some $\{A, B\} \in S_2$ such that $A \cap B = \{x\}$ and neither A nor B equals V . Then $(\{x\}, V) < (A, B)$ and $(\{x\}, V) < (B, A)$, so the consistency of O implies that $(\{x\}, V) \in O$.

Therefore, if $(V, \{x\}) \in O$ then x is not a cutvertex of G . Then, for every other separation $\{A, B\}$, either $x \in A \setminus B$ or $x \in B \setminus A$, and so either $(B, A) < (V, \{x\})$ or $(A, B) < (V, \{x\})$. The consistency of O then determines that

$$O = O_x := \{(A, B) \in \vec{S}_2 : x \in B \text{ and } (A, B) \neq (\{x\}, V)\},$$

which is indeed a profile.

We have shown that every graph contains, for each of its vertices x that is not a cutvertex, a unique 2-profile O_x that is not strongly consistent. However, the orientation

$$O'_x := \{(A, B) \in \vec{S}_2 : x \in B \text{ and } (A, B) \neq (V, \{x\})\}$$

of S_2 is also a 2-profile which does contain every small separation in \vec{S}_2 . (Indeed, $O'_x = O(b)$ for the unique block b containing x .) Since every graph contains a vertex which is not a cutvertex, it follows that

Lemma 9. *Every graph G contains a regular 2-profile.* \square

Lemma 9 means that our goal to find a duality theorem for k -profiles in graphs has substance only for $k > 2$, for which Lemma 8 tells that all k -profiles are regular. In our pursuit of Theorems 2 and 3 it will therefore suffice to study regular \mathcal{F} -tangles of submodular separation systems \vec{S} , such as \vec{S}_k for $\mathcal{F} = \mathcal{P}$.

So, until further notice:

Let \vec{S} be any submodular separation system in some universe \vec{U} , and let \mathcal{F} be a subset of $2^{\vec{S}}$ containing $\mathcal{P} \cap 2^{\vec{S}}$.

Our aim will be to prove a duality theorem for the regular \mathcal{F} -tangles of S .

It will be instructive to keep in mind, as an example, the case of k -blocks, where $\mathcal{F} = \mathcal{B}_k$. In this case any triple $\{(A, B), (C, D), (D \cap B, A \cup C)\} \in \mathcal{P} \cap 2^{\vec{S}_k}$ is contained in \mathcal{B}_k , as $|B \cap D \cap (A \cup C)| < k$ since $(D \cap B, A \cup C) \in S_k$.

For ease of notation, let us write $\mathcal{P}_S := \mathcal{P} \cap 2^{\vec{S}}$, and put $\mathcal{P}_k := \mathcal{P}_{S_k}$ when U is the set of separations of a given graph. Note that an orientation of S avoids \mathcal{P} if and only if it avoids \mathcal{P}_S , and an S -tree is over \mathcal{P} if and only if it is over \mathcal{P}_S .

Our first problem is that, in order to apply Theorem 5, we need \mathcal{F} to be a set of stars. Since our assumptions about \mathcal{F} do not require this, our first aim is to turn \mathcal{F} into a set \mathcal{F}^* of stars such that the regular \mathcal{F} -tangles of S are precisely its regular \mathcal{F}^* -tangles.

Suppose we have some pair of separations \vec{x}_1 and \vec{x}_2 which are both contained in some set $\sigma \subseteq \vec{S}$. Since S , by assumption, is submodular, at least one of $\vec{x}_1 \wedge \vec{x}_2$ and $\vec{x}_2 \wedge \vec{x}_1$ must also be in \vec{S} . To *uncross* \vec{x}_1 and \vec{x}_2 in σ we replace $\{\vec{x}_1, \vec{x}_2\}$ with the pair $\{\vec{x}_1 \wedge \vec{x}_2, \vec{x}_2\}$ in the first case and $\{\vec{x}_1, \vec{x}_2 \wedge \vec{x}_1\}$ in the second case. We note that, in both cases the new pair forms a star and is pointwise \leq the old pair $\{\vec{x}_1, \vec{x}_2\}$. Uncrossing every pair of separations in σ in turn, we can thus turn σ into a star σ^* of separations in at most $\binom{|\sigma|}{2}$ steps, since any star of two separations remains a star if one of its elements is replaced by a smaller separation, and a set of oriented separations is a star as soon as all its 2-subsets are stars. Note, however, that σ^* will not in general be unique, but will depend on the order in which we uncross the pair of separations in σ . Let us say that \mathcal{F}^* is an *uncrossing* of a set \mathcal{F} of sets $\sigma \subseteq \vec{S}$ if

- Every $\tau \in \mathcal{F}^*$ can be obtained by uncrossing a set $\sigma \in \mathcal{F}$;
- For every $\sigma \in \mathcal{F}$ there is some $\tau \in \mathcal{F}^*$ that can be obtained by uncrossing σ .

Note that \mathcal{F}^* , like \mathcal{F} , is a subset of $2^{\vec{S}}$. Also, \mathcal{F}^* contains all the stars from \mathcal{F} , since these have no uncrossings other than themselves. In particular, if \mathcal{F} is standard, i.e. contains all the singleton stars $\{\vec{r}\}$ with \vec{r} trivial in \vec{S} , then so is \mathcal{F}^* .

We have shown the following:

Lemma 10. *\mathcal{F} has an uncrossing \mathcal{F}^* . If \mathcal{F} is standard, then so is \mathcal{F}^* .*

The smaller we can take \mathcal{F}^* to be, the smaller will be the class of S -trees over \mathcal{F}^* . However, to make \mathcal{F}^* as small as possible we would have to give it exactly one star τ for each $\sigma \in \mathcal{F}$, which would involve making a non-canonical choice with regards to the order in which we uncross σ , and possibly which of the two potential uncrossings of a given pair of separations we select.

If we wish for a more canonical choice of family, we can take \mathcal{F}^* to consist of every star that can be obtained by uncrossing a set in \mathcal{F} in any order. Obviously, this will come at the expense of increasing the class of S -trees over \mathcal{F} , i.e. the class of dual objects in our desired duality theorem.

Lemma 11. *Let \mathcal{F}^* be an uncrossing of \mathcal{F} . Then an orientation O of S is a regular \mathcal{F} -tangle if and only if it is a regular \mathcal{F}^* -tangle.*

Proof. Let us first show that if O is a regular \mathcal{F}^* -tangle then it is a regular \mathcal{F} -tangle. It is sufficient to show that O avoids \mathcal{F} . Suppose for a contradiction that there is some $\sigma = \{\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n\} \in \mathcal{F}$ such that $\sigma \subseteq O$. Since \mathcal{F}^* is an uncrossing of \mathcal{F} there is some $\tau = \{\vec{u}_1, \vec{u}_2, \dots, \vec{u}_n\} \in \mathcal{F}^*$ that is an uncrossing of σ . Then, $\vec{u}_i \leq \vec{x}_i \in O$ for all i . Since O is strongly consistent, by Lemma 7, this implies $\vec{u}_i \in O$ for each i . Therefore $\tau \subseteq O$, contradicting the fact that O avoids \mathcal{F}^* .

Conversely suppose O is a regular \mathcal{F} -tangle. We would like to show that O avoids \mathcal{F}^* . To do so, we will show that, if O avoids some set $\sigma = \{\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n\}$ then it also avoids the set $\sigma' = \{\vec{x}_1 \wedge \vec{x}_2, \vec{x}_2 \wedge \vec{x}_3, \dots, \vec{x}_{n-1} \wedge \vec{x}_n\}$ obtained by uncrossing the pair \vec{x}_1, \vec{x}_2 . Then by induction O must also avoid every star obtained by uncrossing a set in \mathcal{F} , and thus will avoid \mathcal{F}^* .

Suppose then that O avoids σ but $\sigma' \subseteq O$. Since $x_1 \in S$ either \vec{x}_1 or \overleftarrow{x}_1 lies in O . As $\sigma \setminus \{\vec{x}_1\} = \{\vec{x}_2, \vec{x}_3, \dots, \vec{x}_n\} \subseteq \sigma' \subseteq O$, but O avoids σ , we have $\vec{x}_1 \notin O$ and hence $\overleftarrow{x}_1 \in O$. But then O contains the triple $\{\overleftarrow{x}_1, \vec{x}_2, \overleftarrow{x}_1 \wedge \vec{x}_2\} \in \mathcal{P}_S \subseteq \mathcal{F}$. This contradicts the fact that O avoids \mathcal{F} . \square

Lemma 11 has an interesting corollary. Suppose, that, in a graph, every star of separations in some given consistent orientation O of S_k points to some k -block. Is there *one* k -block to which all these stars – and hence every separation in O – point? This is indeed the case:

Corollary 12. *If every star of separations in some strongly consistent orientation of S is contained in some profile of S , then there exists one profile of S that contains all these stars.*

Proof. In Lemma 11, take $\mathcal{F} := \mathcal{P}_S$. An orientation O of S whose stars each lie in a profile of S cannot contain a star from \mathcal{P}_S^* . But if O is regular, consistent, and avoids \mathcal{P}_S^* , then by Lemma 11 it also avoids \mathcal{P}_S and hence is a \mathcal{P}_S -tangle. \square

Before we can apply Theorem 5 to our newly found set \mathcal{F}^* of stars, we have to overcome another problem: \vec{S} may fail to be \mathcal{F}^* -separable. To address this problem, let us briefly recall what it means for a family to be closed under shifting. Suppose we have a pair of separations $\vec{r} \leq \vec{s}$ such that \vec{r} is nontrivial, nondegenerate, and not forced by \mathcal{F} . Suppose further that \vec{s} emulates \vec{r} in \vec{S} , and that we have a star $\tau = \{\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n\} \subseteq \vec{S}_{\geq \vec{r}} \setminus \{\vec{r}\}$ that contains some separation $\vec{x}_1 \geq \vec{r}$. Then the image τ' of τ under $f \downarrow_{\vec{s}}^{\vec{r}}$ is

$$\tau' = \{\vec{x}_1 \vee \vec{s}, \vec{x}_2 \wedge \vec{s}, \dots, \vec{x}_n \wedge \vec{s}\},$$

where the fact that \vec{s} emulates \vec{r} guarantees that $\tau' \subseteq \vec{S}$. Let us call τ' a *shift of τ* , and more specifically the *shift of τ from \vec{r} to \vec{s}* . (See [9] for why this is well defined.)

For a family \mathcal{F} to be closed under shifting it is sufficient that it contains all shifts of its elements: that for every $\tau \in \mathcal{F}$, every $\overrightarrow{x_1} \in \tau$, every nontrivial and nondegenerate $\overrightarrow{r} \leq \overrightarrow{x_1}$ not forced by \mathcal{F} , and every \overrightarrow{s} emulating \overrightarrow{r} , the shift of τ from \overrightarrow{r} to \overrightarrow{s} is in \mathcal{F} .

The idea for making Theorem 5 applicable to \mathcal{F}^* will be to close \mathcal{F}^* by adding any missing shifts. Let us define a family $\hat{\mathcal{F}}^*$ as follows: Let $\mathcal{G}_0 = \mathcal{F}^*$, define \mathcal{G}_{n+1} inductively as the set of shifts of elements of \mathcal{G}_n , and put $\hat{\mathcal{F}}^* := \bigcup_n \mathcal{G}_n$. Clearly $\hat{\mathcal{F}}^*$ is closed under shifting.

Next, let us show that a strongly consistent orientation of S avoids \mathcal{F}^* if and only if it avoids $\hat{\mathcal{F}}^*$. We first note the following lemma.

Lemma 13. *Let O be a regular \mathcal{P} -tangle of S . Let $\sigma \subseteq \overrightarrow{S}$ be a star, and let σ' be a shift of σ from some \overrightarrow{r} to some $\overrightarrow{s} \in \overrightarrow{S}$. Then $\sigma' \subseteq O$ implies that $\sigma \subseteq O$.*

Proof. Let $\sigma = \{\overrightarrow{x_1}, \overrightarrow{x_2}, \dots, \overrightarrow{x_n}\}$, with $\overrightarrow{r} \leq \overrightarrow{x_1}$. Then $\overrightarrow{x_1} \vee \overrightarrow{s} \in \sigma' \subseteq O$. Since O is strongly consistent, this implies that $\overrightarrow{x_1}$ and \overrightarrow{s} lie in O . Also, for any $i \geq 2$, as $x_i \in S$, either $\overrightarrow{x_i}$ or $\overleftarrow{x_i}$ lies in O . However, since $\overrightarrow{s} \in O$ and $\overrightarrow{x_i} \wedge \overleftarrow{s} \in \sigma' \subseteq O$, and O avoids \mathcal{P} , it cannot be the case that $\overleftarrow{x_i} \in O$. Hence $\overrightarrow{x_i} \in O$ for all $i \geq 2$, and so $\sigma \subseteq O$. \square

Lemma 14. *Let \mathcal{F}^* be an uncrossing of \mathcal{F} . Then an orientation O of S is a regular $\hat{\mathcal{F}}^*$ -tangle if and only if it is a regular \mathcal{F} -tangle.*

Proof. Recall that O is a regular \mathcal{F} -tangle if and only if it is a regular \mathcal{F}^* tangle (Lemma 11). Clearly every regular $\hat{\mathcal{F}}^*$ -tangle also avoids \mathcal{F}^* , and hence is also a regular \mathcal{F}^* -tangle, and \mathcal{F} -tangle.

Conversely, every regular \mathcal{F} -tangle avoids both $\mathcal{F}^* = \mathcal{G}_0$ (Lemma 11) and hence, by Lemma 13 and $\mathcal{F} \supseteq \mathcal{P}_S$, also \mathcal{G}_1 . Proceeding inductively we see that O avoids \mathcal{G}_n for each n , and so avoids $\bigcup_n \mathcal{G}_n = \mathcal{F}^*$. Hence O is a regular $\hat{\mathcal{F}}^*$ -tangle. \square

Before we can, at last, apply Theorem 5 to our set $\hat{\mathcal{F}}^*$, we have to make one final adjustment: Theorem 5 requires its set \mathcal{F} of stars to be standard, i.e., to contain all singletons stars $\{\overrightarrow{r}\}$ with \overleftarrow{r} trivial in \overrightarrow{S} . Since trivial separations are small, it will suffice to add to $\hat{\mathcal{F}}^*$ all singleton stars $\{\overrightarrow{x}\}$ such that \overleftarrow{x} is small; we denote the resulting superset of $\hat{\mathcal{F}}^*$ by $\overline{\mathcal{F}}^*$. Then $\overline{\mathcal{F}}^*$ -tangles contain all small separations, so they are precisely the regular $\hat{\mathcal{F}}^*$ -tangles.

Clearly, $\overline{\mathcal{F}}^*$ is a standard set of stars. Also, the shift of any singleton star $\{\overrightarrow{x}\}$ is again a singleton star $\{\overrightarrow{y}\}$ such that $\overrightarrow{x} \leq \overrightarrow{y}$. Moreover, if \overleftarrow{x} is small then so is $\overleftarrow{y} \leq \overleftarrow{x}$, so if $\{\overrightarrow{x}\}$ lies in $\overline{\mathcal{F}}^*$ then so does $\{\overrightarrow{y}\}$. Therefore, since $\hat{\mathcal{F}}^*$ is closed under shifting, $\overline{\mathcal{F}}^*$ too is closed under shifting. Hence, we get the following duality theorem for abstract profiles:

Theorem 15. *Let \overrightarrow{S} be a separable⁴ submodular separation system in some universe of separations, let $\mathcal{F} \subseteq 2^{\overrightarrow{S}}$ contain \mathcal{P}_S , and let \mathcal{F}^* be any uncrossing of \mathcal{F} . Then the following are equivalent:*

⁴Whilst the assumption that \overrightarrow{S} is separable is necessary to apply Theorem 5, in a forthcoming paper [6] the authors and Weißauer show that every submodular separation system is in fact separable, and so this assumption can be removed from Theorem 15.

- *There is no regular \mathcal{F} -tangle of S ;*
- *There is no $\overline{\mathcal{F}^*}$ -tangle of S ;*
- *There is an S -tree over $\overline{\mathcal{F}^*}$.*

Proof. By Lemmas 11 to 14 the regular \mathcal{F} -tangles of S are precisely its regular $\overline{\mathcal{F}^*}$ -tangles, and by Lemma 7 these are precisely its $\overline{\mathcal{F}^*}$ -tangles. Hence the first two statements are equivalent. Since $\overline{\mathcal{F}^*}$ is a standard set of stars which is closed under shifting, Theorem 5 implies that the second two statements are equivalent. \square

4 Duality for special tangles, blocks, and profiles

Let us now apply Theorem 15 to prove Theorems 1–3 from the Introduction. We shall first state the latter two results by specifying \mathcal{F} , and then deduce their tree-decomposition formulations as in Theorems 2 and 3, along with Theorem 1. Recall that for any graph the set \overrightarrow{S}_k of separations of order $< k$ is a separable submodular separation system (see [10]).

Theorem 16. *For every finite graph G and every integer $k > 0$ exactly one of the following statements holds:*

- *G contains a k -block;*
- *G has an S_k -tree over $\overline{\mathcal{B}_k^*}$, where \mathcal{B}_k^* is any uncrossing of \mathcal{B}_k .*

Proof. In Theorem 15, let $S = S_k$ be the set of separations of order $< k$ in G , and let $\mathcal{F} := \mathcal{B}_k \cap 2^{\overrightarrow{S}_k}$. Then \overrightarrow{S}_k is a separable submodular separation system in the universe of all separations of G , and the regular \mathcal{F} -tangles in G are precisely the orientations $O(b)$ of S_k for k -blocks b in G . Hence G has a k -block if and only if it has a regular \mathcal{F} -tangle for this \mathcal{F} . The assertion now follows from Theorem 15. \square

As for profiles, every graph G has a regular 1-profile – just orient every 0-separation towards some fixed component – and a regular 2-profile (Lemma 9). So we need a duality theorem only for $k > 2$. Recall that, for $k > 2$, all k -profiles of graphs are regular (Lemma 8).

Theorem 17. *For every finite graph G and every integer $k > 2$ exactly one of the following statements holds:*

- *G has a k -profile;*
- *G has an S_k -tree over $\overline{\mathcal{P}_k^*}$, where \mathcal{P}_k^* is any uncrossing of \mathcal{P}_k .*

Proof. In Theorem 15, let $S = S_k$ be the set of separations of order $< k$ in G , and let $\mathcal{F} := \mathcal{P}_k$. Then \overrightarrow{S}_k is a separable submodular separation system in the universe of all separations of G , and the regular \mathcal{F} -tangles in G are precisely its k -profiles. The assertion now follows from Theorem 15, as before. \square

Theorems 1–3 now follow easily: we just have to translate S -trees into tree-decompositions.

Proof of Theorems 1–3. Given a set S of separations of G and an S -tree (T, α) , with $\alpha(\vec{e}) =: (A_\alpha(\vec{e}), B_\alpha(\vec{e}))$ say, we obtain a tree-decomposition (T, \mathcal{V}_α) of G with $\mathcal{V}_\alpha = (V_t)_{t \in T}$ by letting

$$V_t := \bigcap \{ B_\alpha(\vec{e}) \mid \vec{e} = (s, t) \in \overrightarrow{E(T)} \}.$$

Note that (T, α) can be recovered from this tree-decomposition: given just T and $\mathcal{V} = (V_t)_{t \in T}$, we let α map each oriented edge $\vec{e} = (t_1, t_2)$ of T to the oriented separation of G it induces: the separation $(\bigcup_{t \in T_1} V_t, \bigcup_{t \in T_2} V_t)$ where T_i is the component of $T - e$ containing t_i .

Recall that the set \mathcal{T} defined in (2.2) contains \mathcal{P} . Hence so does any $\mathcal{F} \supseteq \mathcal{T}$. For such \mathcal{F} , therefore, every \mathcal{F} -tangle of S_k is a k -profile, and hence is regular by Lemma 8 if $k > 2$. For $\mathcal{F}_k := \mathcal{F} \cap 2\overrightarrow{S}_k$ and

$$\mathcal{T}_{\mathcal{F}}(k) := \{ (T, \mathcal{V}_\alpha) \mid (T, \alpha) \text{ is an } S_k\text{-tree over } \overrightarrow{\mathcal{F}_k} \}$$

we thus obtain Theorem 1 directly from Theorem 15. Similarly, letting

$$\begin{aligned} \mathcal{T}_{\mathcal{B}}(k) &:= \{ (T, \mathcal{V}_\alpha) \mid (T, \alpha) \text{ is an } S_k\text{-tree over } \overrightarrow{\mathcal{B}_k} \} \\ \mathcal{T}_{\mathcal{P}}(k) &:= \{ (T, \mathcal{V}_\alpha) \mid (T, \alpha) \text{ is an } S_k\text{-tree over } \overrightarrow{\mathcal{P}_k} \} \end{aligned}$$

yields Theorems 2 and 3 as corollaries of Theorems 16 and 17. \square

5 Width parameters

In this section we derive some bounds for the block and profile width of a graph that follow easily from our main results combined with those of [9, 10] and [13].

Given a star σ of separations in a graph, let us call the set $\bigcap \{ B \mid (A, B) \in \sigma \}$ the *interior* of σ in G . For example, every star in \mathcal{P}_k^* is of the form

$$\{ (A, B), (B \cap C, A \cup D), (B \cap D, A \cup C) \} \subseteq S_k,$$

and hence every vertex of its interior lies in at least two of the separators $A \cap B$, $(B \cap C) \cap (A \cup D)$ and $(B \cap D) \cap (A \cup C)$. Since all these separations are in S_k , the interior of any star in \mathcal{P}_k^* thus has size at most $3(k-1)/2$.

We can apply this observation to obtain the following upper bound on the profile-width $\text{pw}(G)$ of a graph G in terms of its tree-width $\text{tw}(G)$:

Theorem 18. *For every graph G ,*

$$\text{pw}(G) \leq \text{tw}(G) + 1 \leq \frac{3}{2}\text{pw}(G). \quad (5.1)$$

Proof. For the first inequality, note that $\text{tw}(G) + 1$ is the largest integer k such that G has no tree-decomposition into parts of order $< k$. By the duality theorem for tree-width from [10], having no such tree-decomposition is equivalent to admitting an \mathcal{S}^k -tangle of S_k , where

$$\mathcal{S}^n = \{ \tau \subseteq \overrightarrow{U} : \tau \text{ is a star and } \left| \bigcap \{ B : (A, B) \in \tau \} \right| < n \}$$

and \overrightarrow{U} is the universe of all separations of G . But among the \mathcal{S}^k -tanglements of S_k are all the k -profiles of G . (An easy induction on $|\tau|$ shows that a regular k -profile has no subset $\tau \in \mathcal{S}^k$; cf. Lemma 8 and [7, Prop. 3.4].) Therefore G has no k -profile for $k > \text{tw}(G) + 1$, which Corollary 4 translates into $\text{pw}(G) \leq \text{tw}(G) + 1$.

For the second inequality, recall that if $k := \text{pw}(G)$ then G has a tree-decomposition in $\mathcal{T}_{\mathcal{P}}(k+1)$. The parts of this tree-decomposition are interiors of stars in $\overline{\mathcal{P}_{k+1}^*}$, so they have size at most $3k/2$. This tree-decomposition, therefore, has width at most $(3k/2) - 1$, which thus is an upper bound for $\text{tw}(G)$. \square

We can also relate the profile-width of a graph to its branch-width, as follows. In order to avoid tedious exceptions for small k , let us define the *adjusted branch-width* of a graph G as

$$\text{brw}(G) := \min\{k \mid G \text{ has no } S_{k+1}\text{-tree over } \mathcal{T}^*\}.$$

By [10, Theorem 4.4], this is equivalent to the *tangle number* of G , the greatest k such that G has a k -tangle. For $k \geq 3$ it coincides with the original branch-width of G as defined by Robertson and Seymour [13].⁵

Robertson and Seymour [13] showed that the adjusted branch-width of a graph is related to its tree-width in the same way as we found that the profile-width is:

$$\text{brw}(G) \leq \text{tw}(G) + 1 \leq \frac{3}{2}\text{brw}(G). \quad (5.2)$$

Together, these inequalities imply the following relationship between branch-width and profile-width:

Corollary 19. *For every graph G ,*

$$\text{brw}(G) \leq \text{pw}(G) \leq \frac{3}{2}\text{brw}(G). \quad (5.3)$$

Proof. The first inequality follows from the fact that k -tangles are k -profiles, and that the largest k for which G has a k -tangle equals the adjusted branch-width: thus,

$$\text{brw}(G) = k \leq \pi(G) = \text{pw}(G)$$

by Corollary 4.

For the second inequality, notice that $\text{pw}(G) \stackrel{(5.1)}{\leq} \text{tw}(G) + 1 \stackrel{(5.2)}{\leq} \frac{3}{2}\text{brw}(G)$. \square

Since every k -block defines a k -profile, Corollary 4 implies that the block-width of a graph is a lower bound for its profile-width, and hence by (5.1) also for its tree-width (plus 1). Conversely, however, no function of the tree-width of a graph can be a lower bound for its block-width. Indeed, the tree-width of the $k \times k$ -grid H_k is at least k (see [5]) but H_k contains no 5-block: in every set of ≥ 5 vertices there are two non-adjacent vertices, and the neighbourhood of either vertex is then a set of size 4 which separates the two vertices. Therefore there exist graphs with bounded block-width and arbitrarily high tree-width.

Since large enough tree-width forces both large profile-width (5.1) and large branch-width (5.2), the grid example shows further that making these latter parameters large cannot force the block-width of a graph above 4.

⁵Our *adjusted branch-width* is the dual parameter to the tangle number for all k , while the original branch-width from [13] achieves this only for $k \geq 3$: it deviates from the tangle number for some graphs and $k \leq 2$. See [10, end of Section 4] for a discussion.

References

- [1] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. k -Blocks: a connectivity invariant for graphs. *SIAM J. Discrete Math.*, 28:1876–1891, 2014.
- [2] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. Connectivity and tree structure in finite graphs. *Combinatorica*, 34(1):1–35, 2014.
- [3] 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.
- [4] R. Diestel. Abstract separation systems. To appear in *Order* (2017), DOI 10.1007/s11083-017-9424-5, arXiv:1406.3797v5.
- [5] R. Diestel. *Graph Theory* (5th edition, 2016). Springer-Verlag, 2017. Electronic edition available at <http://diestel-graph-theory.com/>.
- [6] R. Diestel, J. Erde, and D. Weißauer. Tangles in abstract separation systems. In preparation.
- [7] R. Diestel, F. Hundertmark, and S. Lemanczyk. Profiles of separations: in graphs, matroids, and beyond. arXiv:1110.6207, to appear in *Combinatorica*.
- [8] R. Diestel and S. Oum. Unifying duality theorems for width parameters, II. General duality. arXiv:1406.3798, 2014.
- [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] W. Mader. Über n -fach zusammenhängende Eckenmengen in Graphen. *J. Combin. Theory (Series B)*, 25:74–93, 1978.
- [13] N. Robertson and P.D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)*, 52:153–190, 1991.