

Profiles of separations: in graphs, matroids, and beyond*

Reinhard Diestel Fabian Hundertmark
Sahar Lemanczyk

Mathematisches Seminar, Universität Hamburg

October 19, 2020

Abstract

We show that all the tangles in a finite graph or matroid can be distinguished by a single tree-decomposition that is invariant under the automorphisms of the graph or matroid. This comes as a corollary of a similar decomposition theorem for more general combinatorial structures, which has further applications.

These include a new approach to cluster analysis and image segmentation. As another illustration for the abstract theorem, we show that applying it to edge-tangles yields the Gomory-Hu theorem.

1 Introduction

One of the oldest, and least precise, problems in the connectivity theory of graphs is how to decompose a $(k - 1)$ -connected graph into something like its ‘ k -connected pieces’. The block-cutvertex tree of a graph achieves this for $k = 2$, and Tutte [21] showed for $k = 3$ that every 2-connected finite graph has a tree-decomposition of adhesion 2 whose torsos are either 3-connected or cycles. Tutte’s theorem has recently been extended by Grohe [17] to $k = 4$, and by Carmesin, Diestel, Hundertmark and Stein [4] to arbitrary k as follows.

A k -block of a graph G , where k is any positive integer, is a maximal set X of at least k vertices such that no two vertices $x, y \in X$ can be separated in G by fewer than k vertices other than x and y . We refer to k -blocks for unspecified k simply as *blocks*.

Tutte’s theorem is essentially the case $k = 3$ of the following more recent result. Note however that Theorem 1 does not require the graph to be $(k - 1)$ -connected.

*Cite as in *Combinatorica* **39** (2019), 37–75, and arXiv:1110.6207.

Theorem 1. [4, Theorem 1] *Every finite graph has a canonical tree-decomposition of adhesion $< k$ that efficiently distinguishes all its k -blocks, for every $k \in \mathbb{N}$.*

Here, a tree-decomposition (T, \mathcal{V}) of a graph G with $\mathcal{V} = (V_t)_{t \in T}$ is *canonical* if the automorphisms of G act naturally on T : if they map every part V_t to another part $V_{t'}$ in such a way that $t \mapsto t'$ is an automorphism of T . It *distinguishes* two sets b, b' of vertices in G if T has an edge e whose associated separation of G can be written as $\{A, B\}$ with $b \subseteq A$ and $b' \subseteq B$. (The edge $e = t_1 t_2 \in T$ is *associated with* the separation $\{\bigcup_{t \in T_1} V_t, \bigcup_{t \in T_2} V_t\}$ of G where T_i is the component of $T - e$ containing t_i , and the least order¹ of such a separation is the *adhesion* of (T, \mathcal{V}) .) It distinguishes b and b' *efficiently* if e can be chosen so that $\{A, B\}$ has the least order ℓ of any separation $\{C, D\}$ of G such that $b \subseteq C$ and $b' \subseteq D$. It is not hard to show that $\ell < k$ if b and b' are k -blocks.

More generally, given a k -block b and a $(<k)$ -separation $\{A, B\}$ of G , there cannot be vertices x, y in b such that $x \in A \setminus B$ and $y \in B \setminus A$, as these vertices would then be separated by the set $A \cap B$ of fewer than k vertices. So b will be a subset of A or of B , but not both. In this way, b *orients* $\{A, B\}$: as (A, B) if $b \subseteq B$, and as (B, A) if $b \subseteq A$. (We think of an *oriented separation* (A, B) as ‘pointing towards’ B .)

Every k -block b of G thus induces an *orientation* of the set S_k of all the $(<k)$ -separations of G : a set consisting of exactly one of the two orientations of every separation in S_k . We call this orientation of S_k the *k -profile* of b .

An analysis of the proof of Theorem 1 shows that all it ever uses about k -blocks is their k -profiles. And the theorem itself can be expressed more easily in these terms too. Indeed, it finds a canonical tree-decomposition (T, \mathcal{V}) of the given graph that *distinguishes* the k -profiles of its k -blocks in the following natural sense: for every two such profiles P and P' there is a separation $\{A, B\} \in S_k$ associated with an edge of T that has orientations $(A, B) \in P$ and $(B, A) \in P'$.

The k -profiles of the k -blocks of a graph are not unlike its k -tangles: these, too, are orientations of S_k .² Indeed, the following generalization of Theorem 1 is a corollary merely of rewriting its proof in terms of profiles:

Theorem 2. [2, Theorem 4.5] *Every finite graph has a canonical tree-decomposition of adhesion $< k$ that efficiently distinguishes all its k -blocks and k -tangles, for every fixed $k \in \mathbb{N}$.*

Robertson and Seymour [20] proved that every finite graph has a tree-decomposition that distinguishes all its distinguishable tangles, not just all those of some fixed order k . One of our main results is that we obtain a canonical such decomposition. Using profiles as a common generalization of blocks and tangles,

¹The *order* of a separation $\{A, B\}$ of a graph is the number $|A \cap B|$. Separations of order k are *k -separations*; separations of order $< k$ are *$(< k)$ -separations*.

²Indeed, in Section 2 we shall introduce more abstract ‘ k -profiles’, not necessarily induced by k -blocks, of which k -tangles – those of order k , in the terminology of [20] – are a prime example. See [1, 2] for more on the relationship between profiles, blocks and tangles.

we can even distinguish both at the same time, and from each other, by the same decomposition:

Theorem 3. *Every finite graph has a canonical tree-decomposition that efficiently distinguishes all its distinguishable tangles and robust³ blocks.*

See Section 4.1 for precise definitions.

Representing the k -blocks of a graph G by their k -profiles represents a shift of paradigm for our initial question about those ‘ k -connected components’: k -profiles can identify highly connected regions of G without being concrete objects, such as subgraphs or blocks, simply by pointing to them. In this new paradigm, we would think of these orientations of S_k as the ‘abstract objects’ we wish to distinguish by some tree-decomposition.

Not all orientations of S_k identify highly connected regions of G , though: they must be consistent at least in the sense that no two oriented separations point away from each other. Profiles of blocks are naturally consistent in this sense, and tangles are too, by definition. Other consistent orientations that identify highly connected regions in a graph are explored in [8, 9].

Another advantage of this new paradigm is that we can study, and seek to separate in a tree-like way, highly connected substructures in combinatorial structures other than graphs: all we need is a sensible notion of separation. Tangles for matroids are an example of this, and we obtain the following analogue of Theorem 3:

Theorem 4. *Every finite matroid has a canonical tree-decomposition which efficiently distinguishes all its distinguishable tangles.*

See Section 4.2 for definitions.

Theorem 4 generalizes the matroid analogue of the theorem of Robertson and Seymour mentioned earlier: Geelen, Gerards and Whittle [15] proved that all the distinguishable tangles of a matroid can be distinguished by a single tree-decomposition. But while their tree-decompositions are again not canonical, ours are.

One main aspect of this paper is that tangle-like structures are meaningful, and can be canonically separated in a tree-like way, in a much more general context even than graphs and matroids combined. Just as, for proving our theorems, we do not need to know more about k -blocks than how they orient the separations in S_k , all we need to know about these separations is how they and their orientations are, or fail to be, nested and consistent with each other. This information, however, can be captured by a simple poset with an order-reversing involution defined on the set of these separations: we do not need that they ‘separate’ anything, such as the ground set of a matroid or the vertex set of a graph, into two sides.

³This is a necessary but unimportant restriction; all interesting blocks are robust (Section 3.2).

We shall prove our theorems at this general level of *abstract separation systems* [6]: we define *profiles* as consistent ways of orienting these abstract separations, and find in any separation system S a nested subset T of separations that *distinguish* all the profiles of S , in that for every pair of distinct profiles there exists a separation in T which they orient differently – just as the profiles of two blocks would if these lay on either side of a separation which they both orient.

The fact that we our main theorems can be proved in this abstract setting makes them applicable beyond graphs and matroids. For example, our *canonical tangle-tree theorem*, Theorem 3.6, of which Theorems 3 and 4 will be corollaries, can also be applied to image segmentation and cluster analysis in big data sets. In that context, clusters are described as profiles of bespoke separation systems for the data set in question, which may have nothing to do with graphs or matroids.

This approach to cluster analysis is new. It takes advantage of the fact that profiles, like real-world clusters, can be ‘fuzzy’. For example, consider a large grid in a graph. For every low-order separation, most of the grid will lie on the same side, so the grid ‘orients’ that separation towards this side. But every single vertex will lie on the ‘wrong’ side for *some* low-order separation, the side not containing most of the grid; for example, it may be separated off by its four neighbours. The grid, therefore, defines a unique k -profile for some large k , but the ‘location’ of this profile is not represented correctly by any one of its vertices – just as for a fuzzy cluster of a data set it may be impossible to say which data exactly belong to that cluster and which do not.

Profiles of abstract separation systems can capture such clusters, and our theorems can be applied to describe their relative positions, as soon as the data set comes with a submodular ‘order’ function on its separations. In practice, it appears that most natural ways of cutting a data set in two allow for such order functions, and their choice can be used to specify the exact type of cluster to be analysed. An example from image segmentation is described in [11].

The application of our main result to cluster analysis in large data sets will read as follows:

Theorem 5. *Every submodular data set has a canonical regular tree set of separations which efficiently distinguishes all its distinguishable clusters.*

Abstract separation systems were first introduced in [6]. In Section 2 we give a summary of the basic formal setup, and define profiles as their orientations: a choice of one element from each pair identified by the involution. As our main results, we then prove in Section 3 two general tangle-tree theorems for these abstract separation systems: one canonical, the other not but best possible in other ways. In Section 4 we then apply the first of these to graphs and matroids, obtaining Theorem 3 and Theorem 4 as corollaries.

In Section 5 we first show how to set up abstract separation systems for cluster analysis in large data sets, and show how Theorem 5 follows from our canonical tangle-tree theorem. We then illustrate the non-canonical tangle-tree

theorem by applying it to the ‘edge-tangles’ of a graph. This yields the existence of Gomory-Hu trees known from optimization, reproving the well known Gomory-Hu theorem.

There is also a width-duality theorem for profiles in graphs and matroids, like those for tangles known from [20]: if a given graph, say, has no k -profile for some given k , or no k -profile induced by a k -block, it admits a tree-decomposition witnessing this as an easily checked certificate. This duality theorem will be proved elsewhere [8] but is indicated briefly in Section 6.

We conclude in Section 7 with an open problem for tangles.

Any graph-theoretic terms not defined here can be found in [5]. For matroid terminology we refer to [19]. More on general abstract separation systems can be found in [6], and more on tree sets in [7]. This paper includes and supersedes [18].

2 Abstract separation systems and their profiles

2.1 Separations of sets

Separations in graphs and matroids are certain unordered pairs of subsets of a set: of the vertex set of the graph or the ground set of the matroid. In the case of graphs, the two subsets may overlap; in the case of a matroid, they partition the ground set. But in either case their union is the original set.

Given an arbitrary set V , a *separation* of V is a set $\{A, B\}$ of two subsets A, B such that $A \cup B = V$. Its *order* is the cardinality $|A \cap B|$. Every such separation $\{A, B\}$ has two *orientations*: (A, B) and (B, A) . Inverting these is an involution $(A, B) \mapsto (B, A)$ on the set of these *oriented separations* of V .

The oriented separations of a set are partially ordered as

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

Our earlier involution reverses this ordering:

$$(A, B) \leq (C, D) \Leftrightarrow (B, A) \geq (D, C).$$

The oriented separations of a set form a lattice under this partial ordering, in which $(A \cap C, B \cup D)$ is the infimum of (A, B) and (C, D) , and $(A \cup C, B \cap D)$ is their supremum. The infima and suprema of two separations of a graph or matroid are again separations of that graph or matroid, so these too form lattices under \leq .

Their induced posets of all the oriented separations of order $< k$ for some fixed k , however, need not form such a lattice: when (A, B) and (C, D) have order $< k$, this need not be the case for $(A \cap C, B \cup D)$ and $(A \cup C, B \cap D)$.

2.2 Abstract separation systems

A *separation system* $(\vec{S}, \leq, *)$ is a partially ordered set \vec{S} with an order-reversing involution $*$. An *isomorphism* between two separation systems is a bijection

between their underlying sets that respects their partial orderings and commutes with their involutions.

The elements of a separation system \vec{S} 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 $\vec{r}, \vec{s} \in \vec{S}$,

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

For subsets $R \subseteq \vec{S}$ we write $R^* := \{ \bar{r} \mid \vec{r} \in R \}$.

An (unoriented) *separation* is a set of the form $\{ \vec{s}, \bar{s} \}$, and then denoted by s .⁴ 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} .⁵ Given a set $R \subseteq S$ of separations, we write $\vec{R} := \bigcup R \subseteq \vec{S}$ 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.

A separation $\vec{r} \in \vec{S}$ is *trivial in \vec{S}* , and \bar{r} is *co-trivial*, if there exists $s \in S$ such that $\vec{r} < \vec{s}$ as well as $\vec{r} < \bar{s}$. We call such an s a *witness* of \vec{r} and its triviality. If neither orientation of r is trivial, we call r *nontrivial*. Note that if \vec{r} is trivial in \vec{S} then so is every $\vec{r}' \leq \vec{r}$.

A separation \vec{s} is *small* if $\vec{s} \leq \bar{s}$. Trivial separations are small by (1), but other separations can be small too. But if \vec{s} is small and $\vec{r} < \vec{s}$, then \vec{r} is clearly trivial. So all but the largest small separations are in fact trivial. An unoriented separation is *proper* if it has no small orientation. A separation system is *regular* if none of its elements is small [6].

If there are binary operations \vee and \wedge on a separation system $(\vec{U}, \leq, *)$ 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{U} , we call $(\vec{U}, \leq, *, \vee, \wedge)$ 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 universe \vec{U} is *submodular* if it comes with a submodular *order function*, a real function $\vec{s} \mapsto |\vec{s}|$ on \vec{U} that satisfies $0 \leq |\vec{s}| = |\bar{s}|$ and

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

for all $\vec{r}, \vec{s} \in \vec{U}$. We call $|s| := |\vec{s}|$ the *order* of s and of \vec{s} . For every integer $k > 0$, then,

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

is a separation system. But \vec{S}_k need not itself be a universe, at least not with respect to the operations \vee and \wedge induced from \vec{U} , since suprema or infima in \vec{U} of elements of \vec{S}_k can lie outside \vec{S}_k .

⁴To smooth the flow of the narrative we usually also refer to oriented separations simply as 'separations' if the context or use of the arrow notation \vec{s} shows that they are oriented.

⁵It 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.

An *isomorphism* between universes of separations is a bijection between their ground sets that respects their orderings and commutes with their involutions and, in the case of submodular universes, their order functions.

Separations of a set V , and their orientations, are clearly an instance of this abstract setup if we identify $\{A, B\}$ with $\{(A, B), (B, A)\}$. The small separations of V are those of the form (A, V) , the trivial ones are those of the form (U, V) with $U \subseteq A \cap B$ for some other separation $\{A, B\} \neq \{U, V\}$, and the proper separations are those of the form $\{A, B\}$ with $A \setminus B$ and $B \setminus A$ both nonempty. The separations of V form a submodular universe for the order function $|(A, B)| := |A, B| := |A \cap B|$.

2.3 Tree sets of separations

Given a separation system $(\vec{S}, \leq, *)$, two separations $r, s \in S$ 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.⁶ 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. Two sets of separations are *nested* if every element of the first set is nested with every element of the second. A *tree set* is a nested separation system without trivial or degenerate elements. When $\vec{T} \subseteq \vec{S}$ is a tree set, we also call $T \subseteq S$ a *tree set* (and *regular* if \vec{T} is regular).

For example, the set of orientations (u, v) of the edges uv of a tree T form a regular tree set with respect to the involution $(u, v) \mapsto (v, u)$ and the *natural partial ordering* on $\vec{E}(T)$: the ordering in which $(x, y) < (u, v)$ if $\{x, y\} \neq \{u, v\}$ and the unique $\{x, y\}$ - $\{u, v\}$ path in T joins y to u . The oriented bipartitions of $V(T)$ defined by deleting an edge of T form a tree set isomorphic to this one.

The separations of a graph associated with a tree-decomposition are also nested. But since they may be small, or even trivial, they need not form a tree set, not even an irregular one. But conversely, every tree set of separations of a finite graph G comes from a tree-decomposition of G , which is essentially unique [7].

Any two elements r, s of a universe \vec{S} of separations have four *corner separations*

$$\{(\vec{r} \vee \vec{s}), (\vec{r} \vee \vec{s})^*\}, \{(\vec{r} \vee \vec{s}), (\vec{r} \vee \vec{s})^*\}, \{(\vec{r} \vee \vec{s}), (\vec{r} \vee \vec{s})^*\}, \{(\vec{r} \vee \vec{s}), (\vec{r} \vee \vec{s})^*\}$$

(see Figure 1). By (2), these can also be expressed in terms of \wedge .

Lemma 2.1. *Let $r, s \in \vec{S}$ be two crossing separations. Every separation t that is nested with both r and s is also nested with all four corner separations of r and s .*

⁶Terms introduced for unoriented separations may be used informally for oriented separations too if the meaning is obvious, and vice versa.




Figure 1: Separations as in Lemma 2.1

Proof. Since t is nested with r and s , it has an orientation pointing towards r , and one pointing towards s . If these orientations of t are not the same, then $\vec{r} \leq \vec{t} \leq \vec{s}$ for suitable orientations of r, s, t . In particular, r and s are nested, contrary to our assumption. Hence t has an orientation \vec{t} that points towards both r and s .

Now r and s have orientations $\vec{t} \leq \vec{r}$ and $\vec{t} \leq \vec{s}$. Since \wedge and \vee denote infima and suprema in \vec{S} , we have $\vec{t} \leq \vec{r} \wedge \vec{s} = (\vec{r} \vee \vec{s})^*$ by (2), as well as trivially $\vec{t} \leq \vec{r} \vee \vec{s}$ and $\vec{t} \leq \vec{r} \vee \vec{s}$ and $\vec{t} \leq \vec{r} \vee \vec{s}$. \square

2.4 Profiles of separation systems

Given a separation system $(\vec{S}, \leq, *)$, a subset $O \subseteq \vec{S}$ is an *orientation* of S (and of \vec{S}) if $O \cup O^* = \vec{S}$ and $|O \cap \{\vec{s}, \vec{s}^*\}| = 1$ for all $s \in S$. Thus, O contains every degenerate separation from \vec{S} and contains exactly one orientation of every nondegenerate one. For subsets $S' \subseteq S$ we say that O *induces* and *extends* the orientation $O \cap \vec{S}'$ of S' , and thereby *orients* S' .

A set $O \subseteq \vec{S}$ is *consistent* if there are no distinct $r, s \in S$ with orientations $\vec{r} < \vec{s}$ such that $\vec{r}, \vec{s} \in O$. Consistent orientations of \vec{S} contain all its trivial separations: if \vec{r} is trivial in \vec{S} , witnessed by s say, then s cannot be oriented consistently with \vec{r} . Since s must be oriented somehow, this implies that $\vec{r} \notin O$ and hence $\vec{r} \in O$.

Assume now that $(\vec{S}, \leq, *)$ lies inside a universe $(\vec{U}, \leq, *, \vee, \wedge)$. Generalizing our notion of the profiles of blocks in graphs from the introduction, let us call an orientation P of S a *profile* (of S or \vec{S}) if it is consistent and satisfies

$$\text{For all } \vec{r}, \vec{s} \in P \text{ the separation } \vec{r} \wedge \vec{s} = (\vec{r} \vee \vec{s})^* \text{ is not in } P. \quad (\text{P})$$

Thus if P contains \vec{r} and \vec{s} it also contains $\vec{r} \vee \vec{s}$, unless $\vec{r} \vee \vec{s} \notin \vec{S}$.

A profile is *regular* if it contains no separation whose inverse is small. Examples of irregular profiles are those we call *special*: profiles P of S for which there exists a separation $\vec{s} \in P$ such that \vec{s} is small and

$$P = \{ \vec{r} \in \vec{S} \mid \vec{r} \leq \vec{s} \} \setminus \{ \vec{s} \}.$$

In general, S can have irregular profiles that are not special, but this will not happen in the cases relevant to us.

Applying (P) to any degenerate separation shows that if S has a degenerate element it has no profile.

Note that every subset Q of a profile of \vec{S} is a profile of $\vec{R} = Q \cup Q^* \subseteq \vec{S}$. Put another way, if P is a profile of S and $R \subseteq S$, then $P \cap \vec{R}$ is a profile of R , which we say is *induced* by P . If \mathcal{P} is a set of profiles of S , we write

$$\mathcal{P} \upharpoonright R := \{ P \cap \vec{R} \mid P \in \mathcal{P} \}$$

for the set of profiles they induce on R , and say that \mathcal{P} *induces* $\mathcal{P} \upharpoonright R$.

For example, the k -profile of a k -block b in a graph G , as defined in the introduction, is the orientation $\{(A, B) \in \vec{S}_k \mid b \subseteq B\}$ of the set S_k of separations of G of order $< k$. This orientation of S_k is consistent and satisfies (P), and is thus a (regular) profile of S_k . Similarly, every k -tangle of G is a regular profile of S_k ; see Section 4.1.

When \vec{U} is a universe of separations, we call the profiles of the separation systems $\vec{S}_k \subseteq \vec{U}$ of its separations of order $< k$ the *k -profiles in \vec{U}* . Even if we are interested only in the k -profiles for some fixed k , it is usually important to have \vec{U} , rather than just \vec{S}_k , as a larger universe in which \vee and \wedge are defined. Here is an example:

Example 1. Let G be a graph with a tree-decomposition whose tree is a 3-star. Let the central part in this decomposition be a triangle xyz and the other three parts pendant edges xx' , yy' and zz' . Let \vec{r} , \vec{s} , \vec{t} be the separations splitting G at x , y , z respectively, i.e., $\vec{r} = (\{x', x\}, \{x, y, z, y', z'\})$, and similarly \vec{s} for y and \vec{t} for z . Then $P = \{\vec{r}, \vec{s}, \vec{t}\}$, together with all the small separations of the form (v, V) , is a 2-profile of G in the universe of all its separations, which points to its central triangle. Indeed, P satisfies (P) vacuously because, for example, $\vec{r} \vee \vec{s}$ is a separation of order 2, which does not get oriented by P .

However if the universe in which \vee and \wedge are defined was just the set $S = \{r, s, t\}$ of all the proper 1-separations of G , then $\vec{r} \vee \vec{s} = \vec{t}$ would violate (P). In this sparser universe, the only 2-profiles would be the orientations of S orienting all its separations towards one of the leaves x', y', z' : it would lack the ‘resolution’ for its profiles to see the central triangle as a highly connected substructure.

Let us close this general section on profiles by showing that even in contexts where irregular profiles do occur, they are often just the special ones:

Proposition 2.2. *If \vec{S} contains all trivial separations in \vec{U} that have a witness in S , then every irregular profile P of S is special.*

Proof. Since P is irregular, it contains a separation \vec{s} whose inverse \vec{s} is small. If P is not special, it contains another separation $\vec{r} \not\leq \vec{s}$. Then also $\vec{s} \not\leq \vec{r}$, since otherwise $\vec{r} < \vec{s} \leq \vec{s}$ would be trivial and hence in P . Hence $\vec{r}, \vec{s} < \vec{r} \vee \vec{s}$. But now $(\vec{r} \vee \vec{s})^* < \vec{s} \leq \vec{s}$ is trivial, hence in S by assumption, and thus also in P . But this means that P violates (P), a contradiction. \square

3 Tangle-tree theorems for profiles

As explained in the introduction, the paradigm proposed in this paper is that certain consistent orientations of a separation system \vec{S} , such as profiles, can be treated like highly connected substructures of some structure which these separations separate, even if \vec{S} is given abstractly and no such structure is known. We shall therefore be interested in finding small subsets T of S that ‘distinguish’ some given profiles of S , in the following natural sense.

We say that a non-degenerate separation $s \in S$ *distinguishes* two orientations O, O' of subsets of S if s has orientations $\vec{s} \in O$ and $\bar{s} \in O'$. (We then also say that \vec{s} and \bar{s} themselves distinguish O from O' .) The sets O, O' are *distinguishable* if there exists some $s \in S$ that distinguishes them. A set $T \subseteq S$ *distinguishes* O from O' if some $s \in T$ distinguishes them, and T *distinguishes* a set \mathcal{O} of orientations of subsets of S if it distinguishes its elements pairwise.

A separation r and its orientations are *\mathcal{O} -relevant* if r distinguishes some two elements of \mathcal{O} . A set of separations is *\mathcal{O} -relevant* if all its elements are.

Later, these sets \mathcal{O} will typically be sets \mathcal{P} of profiles: either of profiles of some fixed separation system \vec{S} , or of profiles in some fixed universe \vec{U} of separations. (Note that a separation $r \in U$ can distinguish an ℓ -profile in \vec{U} from a k -profile, for $\ell \neq k$, since both are subsets of \vec{U} .)

No trivial separation in \vec{S} can ever distinguish two profiles of S , because every consistent orientation of S , and in particular every profile, contains it. Since trivial separations in \vec{S}_ℓ are also trivial in any $\vec{S}_k \supseteq \vec{S}_\ell$, this implies that trivial separations in \vec{S}_ℓ cannot distinguish an ℓ -profile from a k -profile in \vec{U} for any $\ell \leq k$.

However, a separation \vec{s} distinguishing two profiles can be small; indeed there can be profiles of the same separation system \vec{S} that differ only in how they orient such a separation s . For $\ell < k$, however, a small separation $\vec{s} \in \vec{S}_\ell$ can be nontrivial in \vec{S}_ℓ but trivial in \vec{S}_k . Such a separation s can distinguish an ℓ -profile (which orients it as \bar{s}) from a k -profile (which must orient it as \vec{s}) in a common universe \vec{U} .

3.1 Tree sets distinguishing profiles of a separation system

Let $(\vec{S}, \leq, *)$ be a separation system insided some universe \vec{U} of separations, and \mathcal{P} a set of profiles of S .

Distinct profiles of S are always distinguished by some $s \in S$, so S distinguishes \mathcal{P} . Our aim is to find a nested (and, in particular, small) subset T of S that still distinguishes all of \mathcal{P} . Since trivial separations do not help in distinguishing profiles, \vec{T} will normally be a tree set, though not necessarily a regular one.

Moreover, we shall often ask that T should be *canonical* for S : that the map $(\vec{S}, \mathcal{P}) \mapsto \vec{T}$ should commute with isomorphisms of separation systems. In particular, \vec{T} will then be invariant (globally) under the automorphisms of \vec{S} if \mathcal{P} is.

Let us start with an example showing that our task of finding a canonical nested subset $T \subseteq S$ that still distinguishes all of \mathcal{P} is nontrivial, indeed will

not be achievable without further assumptions on \mathcal{P} (such as (3) below):

Example 2. Suppose S consists of two crossing separations r and s , and \mathcal{P} consists of the four profiles $\{\vec{r}, \vec{s}\}, \{\vec{r}, \vec{s}\}, \{\overleftarrow{r}, \vec{s}\}, \{\overleftarrow{r}, \vec{s}\}$ of S . Since only the singleton subsets of S are nested, but each of r and s distinguishes only one pair of profiles, S has no nested subset T that distinguishes \mathcal{P} .

The problem with this example is, very loosely, that \mathcal{P} is too large and diverse compared with S . Indeed we can mend it by increasing S , without materially changing \mathcal{P} , as follows.

Example 3. Add to \vec{S} the separations $\vec{t} := \vec{r} \vee \vec{s}$ and $\overleftarrow{t} := \vec{r} \wedge \vec{s}$ together with their inverses. Now \mathcal{P} is no longer formally a set of profiles of this larger separation system \vec{S}' , but in substance it is: every $P \in \mathcal{P}$ extends uniquely to a profile of \vec{S}' . (Indeed, by (P) and consistency, $\{\vec{r}, \vec{s}\} \in \mathcal{P}$ extends only to $\{\vec{r}, \vec{s}, \vec{t}, \overleftarrow{t}\}$ while $\{\overleftarrow{r}, \vec{s}\}$ extends only to $\{\overleftarrow{r}, \vec{s}, \overleftarrow{t}, \vec{t}\}$; for the other two elements of \mathcal{P} consistency forces us to add \overleftarrow{t} and \vec{t} . But these extensions are indeed profiles of S' .) And now there are two tree sets in S' that distinguish these four profiles, $\{r, t, t'\}$ and $\{s, t, t'\}$.

Neither of these tree sets is canonical as yet, since they map to each other under an automorphism of \vec{S}' . However if we add the other two corner separations too, our separation system will be rich enough: our four profiles still remain materially unchanged, but the four corner separations form a canonical tree set that distinguishes them all.

This motivates the following condition, which we shall see will be sufficient for the existence of a canonical tree set $T \subseteq S$ that still distinguishes \mathcal{P} . Let us say that S *scatters* \mathcal{P} if the following holds:

$$\begin{aligned} & \text{Whenever } \vec{r}, \vec{s} \in \vec{S} \text{ cross and there are profiles } P, P' \in \mathcal{P} \text{ with} \\ & \vec{r}, \vec{s} \in P \text{ and } \overleftarrow{r}, \vec{s} \in P', \text{ there exists a separation } \vec{t} \in P \text{ such that} \quad (3) \\ & \vec{r} \vee \vec{s} \leq \vec{t}. \end{aligned}$$

(By symmetry, there is then also be a separation $\overleftarrow{t} \in P'$ such that $\overleftarrow{r} \vee \vec{s} \leq \overleftarrow{t}$.)

Note that the separation t in (3) distinguishes P from P' , so t is \mathcal{P} -relevant. Indeed while $\vec{t} \in P$, we have $\overleftarrow{t} \leq \vec{s} \in P'$, so $\overleftarrow{t} \in P'$ by the consistency of P' unless $s = t$. But $s = t$ is impossible, since r crosses s but is nested with t .

We shall also say that a subset $R \subseteq S$ *scatters* \mathcal{P} if R distinguishes \mathcal{P} and scatters the set $\mathcal{P} \upharpoonright R$ of profiles of R that \mathcal{P} induces. For example, if S scatters \mathcal{P} then so does the set R of all \mathcal{P} -relevant separations in S .

Condition (3) will help us find our nested set $T \subseteq S$ that still distinguishes \mathcal{P} , as follows. Suppose we are trying to pick T from S inductively. At some point we might wish to select for T some separation that distinguishes two profiles $P, P' \in \mathcal{P}$, and have two separations $\vec{r}, \vec{s} \in \vec{S}$ that both do this. If \vec{r} and \vec{s} cross, we can put at most one of them in T . Condition (3) will now save us having to choose: we can pick \vec{t} instead, as this also distinguishes P from P' . Moreover, since t is nested with both r and s , adding it to T will allow us to keep our options open about r or s : if they are eligible for T now, they will still

be eligible once we have added t , so we can still settle for either one of them (or neither) later.

An ideal starting point for finding a tree set $T \subseteq S$ of separations that still distinguishes \mathcal{P} would be a separation t that distinguishes some pair of profiles in \mathcal{P} and is nested with *all* of S : putting such a separation t in T will achieve something, and will do no harm to the desired nestedness of our eventual T .

Looking for such a separation t seems like a long shot. And indeed, condition (3) cannot ensure its existence: since (3) speaks only about \mathcal{P} -relevant separations, the separation t it provides need not be nested with all of S . But since only the \mathcal{P} -relevant separations in S matter for us, the following lemma will suffice:

Lemma 3.1. *If S is \mathcal{P} -relevant and scatters \mathcal{P} , then every maximal $\vec{s} \in \vec{S}$ is nested with all the other separations in S .*

Proof. Suppose some $r \in S$ crosses s . Our aim is to find a pair $P, P' \in \mathcal{P}$ and an orientation \vec{r} of r such that $\vec{r}, \vec{s} \in P$ and $\vec{r}, \vec{s} \in P'$: then $\vec{s} < \vec{t}$ for the separation \vec{t} from (3), contradicting the maximality of \vec{s} . (The inequality is strict, because t is nested with r but s is not.)

As s is \mathcal{P} -relevant, there are profiles $P, P' \in \mathcal{P}$ such that $\vec{s} \in P$ and $\vec{s} \in P'$. Let \vec{r} be the orientation of r that lies in P . If $\vec{r} \in P'$ we are done, so assume that also $\vec{r} \in P'$. Since r is \mathcal{P} -relevant, there exists $P'' \in \mathcal{P}$ such that $\vec{r} \in P''$. If P'' contains \vec{s} , renaming P'' as P' yields the desired pair. If P'' contains \vec{s} , renaming P'' as P and swapping the names of \vec{r} and \vec{r} yields the desired pair P, P' . \square

Every \vec{s} as in Lemma 3.1 distinguishes a unique profile $P_{\vec{s}}$ from the rest of \mathcal{P} :

Lemma 3.2. *If S scatters \mathcal{P} , and $\vec{s} \in \vec{S}$ is maximal among the \mathcal{P} -relevant separations in \vec{S} , then \vec{s} lies in a unique profile $P_{\vec{s}} \in \mathcal{P}$. Thus, s distinguishes $P_{\vec{s}}$ from every other profile in \mathcal{P} .*

Proof. As s is \mathcal{P} -relevant, there are profiles $P, P' \in \mathcal{P}$ such that $\vec{s} \in P$ and $\vec{s} \in P'$.

To show that $P_{\vec{s}} := P$ is unique, suppose there is another profile $P'' \in \mathcal{P}$ that contains \vec{s} . As $P'' \neq P$, there exists $r \in S$ with orientations $\vec{r} \in P$ and $\vec{r} \in P''$. Since P and P'' orient s identically but r differently, we have $r \neq s$. Hence if r is nested with s , then $\vec{r} < \vec{s}$ or $\vec{r} < \vec{s}$ by the maximality of \vec{s} . This contradicts the consistency of P'' or P , respectively.

So r crosses s . If $\vec{r} \in P'$, we apply (3) to find $\vec{s} < \vec{t} \in P$, which contradicts the maximality of \vec{s} . If $\vec{r} \in P'$, we apply (3) with P'' (as the P in (3)) and P' , swapping the names of \vec{r} and \vec{r} , to obtain $\vec{s} < \vec{t} \in P''$ with a similar contradiction. \square

Consider an automorphism $\alpha: \vec{s} \mapsto \vec{s}^\alpha$ of the separation system $(\vec{S}, \leq, *)$. For subsets P of \vec{S} we write $P^\alpha := \{ \vec{s}^\alpha \mid \vec{s} \in P \}$, and put $\mathcal{P}^\alpha := \{ P^\alpha \mid P \in \mathcal{P} \}$ for sets of such subsets. Since α preserves (P) and consistency, \mathcal{P}^α is again a set of profiles of S .

Note that α acts naturally also on the set S of unoriented separations, mapping $s = \{\vec{s}, \bar{s}\}$ to $s^\alpha := \{\vec{s}^\alpha, \bar{s}^\alpha\}$. For sets $T \subseteq S$ we write $T^\alpha := \{s^\alpha \mid s \in T\}$. If $R \subseteq S$ scatters \mathcal{P} , then clearly R^α scatters \mathcal{P}^α .

We are now ready to prove that separation systems scattering a set of profiles contain canonical tree sets that still distinguish these profiles. The main trick in the proof, which builds this tree set T inductively, is to toggle between adding separations to T , which decreases the set \mathcal{P} of profiles not yet distinguished, and reducing S to its \mathcal{P} -relevant subset, so that we can re-apply Lemmas 3.1 and 3.2.

Theorem 3.3. (Canonical tangle-tree theorem for separation systems)

*For every finite separation system $(\vec{S}, \leq, *)$ inside some universe of separations, and every set \mathcal{P} of profiles of S which S scatters, there is a tree set $T \subseteq S$ of separations such that*

- (i) T distinguishes \mathcal{P} ;
- (ii) T is \mathcal{P} -relevant;
- (iii) T is regular if all $P \in \mathcal{P}$ are regular, or if every separation in S is \mathcal{P} -relevant and no $P \in \mathcal{P}$ is special.

*These sets $T = T(S, \mathcal{P})$ can be chosen so that $T(S, \mathcal{P})^\alpha = T(S', \mathcal{P}^\alpha)$ for every isomorphism $\alpha: (\vec{S}, \leq, *) \rightarrow (\vec{S}', \leq, *)$ of separation systems. In particular, T is invariant under any automorphism of $(\vec{S}, \leq, *)$ that maps \mathcal{P} to itself.*

Theorem 3.3 will be the basic building block for a more general canonical tangle-tree theorem we shall prove in the next section. There, we shall find tree sets in universes of separations whose profiles they can distinguish even when these are not profiles of the same separation system.

Proof of Theorem 3.3. We begin by proving the theorem without statement (iii), applying induction on $|\mathcal{P}|$ with both S and \mathcal{P} variable. For $|\mathcal{P}| \leq 1$ there is no pair of profiles to distinguish, so $T(S, \mathcal{P}) := \emptyset$ satisfies (i) and (ii) trivially. For the last statement note that $|\mathcal{P}^\alpha| = |\mathcal{P}|$, so $T(S, \mathcal{P})^\alpha = \emptyset = T(S', \mathcal{P}^\alpha)$.

Assume now that $|\mathcal{P}| \geq 2$. Let $R \subseteq S$ be the set of \mathcal{P} -relevant separations in S . As $|\mathcal{P}| \geq 2$ we have $R \neq \emptyset$. As remarked after (3), R still scatters \mathcal{P} . By Lemma 3.1 applied to $\mathcal{P} \upharpoonright R$, the set R^+ of maximal elements of \vec{R} is nested with all of R , and hence so is the set T^+ of all (unoriented) separations in R that have an orientation in R^+ . In particular, T^+ itself is nested.

By Lemma 3.2, each $\vec{s} \in R^+$ distinguishes a unique profile $P_{\vec{s}} \in \mathcal{P} \upharpoonright R$ from the rest of $\mathcal{P} \upharpoonright R$. Let

$$\mathcal{P}^- := \mathcal{P} \upharpoonright R \setminus \{P_{\vec{s}} \mid \vec{s} \in R^+\}.$$

As $|\mathcal{P}^-| < |\mathcal{P}|$, our induction hypothesis provides a nested set $T^- \subseteq R$ that satisfies (i) and (ii) for \mathcal{P}^- . Let

$$T := T^+ \cup T^-.$$

Clearly, T is nested. It is \mathcal{P} -relevant since $T \subseteq R$. Thus, T is a tree set satisfying (ii). By (i) for \mathcal{P}^- and the choice of T^+ , it distinguishes $\mathcal{P} \upharpoonright R$ and hence also \mathcal{P} , so T also satisfies (i).

The last statement of the theorem is straightforward to check. It follows from the fact that if a separation $r \in S$ is \mathcal{P} -relevant then r^α is \mathcal{P}^α -relevant, and that α maps the maximal elements of \vec{R} to those of \vec{R}^α . This completes our proof of the theorem without statement (iii).

For a proof of (iii), suppose T is irregular. Then some $s \in T$ has a small orientation, \vec{s} say. By (ii) we have $s \in R$, so there exists $P \in \mathcal{P}$ with $\vec{s} \in P$. This P is irregular. To complete the proof of (iii), we show that if $S = R$ then P is special.

Note that \vec{s} is maximal in \vec{R} : for any $\vec{r} > \vec{s}$ in \vec{S} its inverse $\vec{r} < \vec{s} \leq \vec{s}$ is trivial, so $r \notin R$, because only nontrivial separations can distinguish profiles. By our assumption that $S = R$ and Lemma 3.1, the separation s is nested with every other separation in S . By its maximality, \vec{s} points to no other separation in S , so every $r \in S \setminus \{s\}$ has an orientation $\vec{r} < \vec{s} \in P$. By its consistency, P contains \vec{r} rather than \vec{r} . Thus, $P = \{ \vec{r} \mid \vec{r} \leq \vec{s} \} \setminus \{ \vec{s} \}$, i.e., P is special. \square

Let us add a few remarks. It may seem that assertion (ii) would come for free once we have a set T satisfying the other assertions, simply by choosing it inclusion-minimal subject to (i). This, however, will lose us the last statement, canonicity. Indeed it can happen that T cannot be chosen minimal:

Example 4. Let G be a graph obtained by extending a complete graph on k vertices to three otherwise disjoint complete subgraphs B_1, B_2, B_3 of $n > k$ vertices each. Each of these B_i is an n -block that induces an n -profile on the set S of the three k -separations $\{A_i, B_i\}$ of G consisting of (A_i, B_i) and (B_j, A_j) for $j \neq i$. For the set \mathcal{P} of these profiles, Theorem 3.3 finds $T = S$. This set is not minimal with assertion (i), since any two of the three separations in T suffice to separate all three profiles. But only if T contains all three of them will it be closed under the automorphisms of \vec{S} , which all map \mathcal{P} to itself.

Note also that without any assumptions about \mathcal{P} , such as the premise in (iii), there need not be a regular tree set T as in Theorem 3.3. Here is a typical example:

Example 5. Let \vec{S} itself be a tree set, inside some universe \vec{U} of separations, such that $\vec{r} \vee \vec{s} \in \vec{U} \setminus \vec{S}$ for all incomparable $\vec{r}, \vec{s} \in \vec{S}$. Let \vec{s} be maximal in \vec{S} . Then every $r \in S$ has an orientation $\vec{r} \leq \vec{s}$, so $P = \{ \vec{r} \mid \vec{r} \leq \vec{s} \}$ is a profile of S . By replacing \vec{s} with \vec{s} in P we obtain another profile P' of S : note that P' is again consistent, since by the maximality of \vec{s} in \vec{S} its new separation \vec{s} cannot be part of an inconsistent pair of separations. Then s distinguishes P from P' , it is the only separation that does, and hence it will lie in the tree set T satisfying Theorem 3.3.

However, it is easy to construct instances of this where \vec{s} is small, making P special. For example, S might be the set of separations induced by a tree-decomposition of a graph $G = (V, E)$, with $\vec{s} = (V, A)$ corresponding to an edge uv at a leaf v of the decomposition tree. In this case T will not be regular, because it contains the improper separation s .

3.2 Tree sets distinguishing profiles in separation universes

Let $\vec{U} = (\vec{U}, \leq, *, \vee, \wedge, | |)$ be a submodular universe of separations. Our aim is to show that \vec{U} contains a canonical tree set T that distinguishes all its profiles.

Theorem 3.3 provides a start for fixed k :

Corollary 3.4. *If \vec{U} contains no separation of order less than $k - 1$, then there is a tree set $T \subseteq S_k$ that distinguishes all the k -profiles in \vec{U} and is invariant under its automorphisms.*

Proof. This follows from Theorem 3.3 once we have checked that S_k scatters all its profiles, the k -profiles in \vec{U} . Note first the separations in S_k have order exactly $k - 1$, by assumption. Given $\vec{r}, \vec{s} \in P$ as in (3), we have $|\vec{r} \wedge \vec{s}| \geq k - 1$ by assumption, and hence $|\vec{r} \vee \vec{s}| \leq k - 1$ by submodularity, since r and s have order exactly $k - 1$. Hence $\vec{t} := \vec{r} \vee \vec{s} \in \vec{S}_k$, and therefore $\vec{t} \in P$ by (P) and $\vec{r}, \vec{s} \in P$. \square

In general we do not have the helpful assumption that $|s| \geq k - 1$ for all $s \in U$. But our plan is to mimic it, by considering for $k = 1, 2, \dots$ in turn those profiles that can be distinguished by a separation of order $< k$. At each step, we may then assume inductively that all the profiles we still need to distinguish cannot be distinguished by a separation of smaller order, just as in the premise of Corollary 3.4.

The challenge with this approach is that, while each application of Theorem 3.3 provides a tree set of separations, we also have to ensure that these tree sets are nested with each other. We shall achieve this in the end, but it will need some care.

Let us say that two profiles in \vec{U} are k -distinguishable if some separation of order at most k distinguishes them. The smallest k for which distinguishable profiles P, P' in \vec{U} are k -distinguishable is denoted by $\kappa(P, P')$. If $s \in U$ distinguishes P from P' and has order $|s| = \kappa(P, P')$, we say that s, \vec{s} and \vec{s} distinguish P from P' *efficiently*. A set $T \subseteq U$ distinguishes a set \mathcal{P} of profiles in \vec{U} *efficiently* if every two profiles in \mathcal{P} are distinguished efficiently by some separation in T .

Consider a pair of crossing separations $\vec{r}, \vec{s} \in \vec{U}$ and their corner separations $\vec{r}_1 := \vec{r} \wedge \vec{s}$ and $\vec{r}_2 := \vec{r} \wedge \vec{s}$. Then $\vec{r}_1, \vec{r}_2 \leq \vec{r}$. If \vec{r} is the supremum in \vec{U} of \vec{r}_1 and \vec{r}_2 , i.e., if $\vec{r} = \vec{r}_1 \vee \vec{r}_2$,⁷ then \vec{r}_1, \vec{r}_2 and \vec{r} cannot lie in a common profile, because this would violate (P). But in general \vec{r}_1, \vec{r}_2 and \vec{r} might lie in a common profile, say in the k -profile P . This can only happen if $|s| \geq k$, since otherwise P would contain \vec{s} or \vec{s} , but both $\{\vec{s}, \vec{r}, \vec{r}_1\} \subseteq P$ and $\{\vec{s}, \vec{r}, \vec{r}_2\} \subseteq P$ would violate (P). If it does happen then, intuitively, the separation $s \in U \setminus S_k$ looks a bit as though it splits P in half. This can cause problems, so let us give a name to certain profiles that are not split in this way.

⁷If U is the set of all bipartitions of a set, then this is always the case. Profiles in matroids, therefore, or clusters in data sets, will always be robust in the sense defined below.

Let us say that a profile P in \vec{U} is *robust* if it is n -robust for every n , and *n-robust* if for every $\vec{r} \in P$ and every $s \in S_n$ the following holds:

$$\text{If } \vec{r} \wedge \vec{s} \text{ and } \vec{r} \wedge \vec{s} \text{ both have order } < |r|, \text{ they do not both lie in } P. \quad (\text{R})$$

Every n -profile is n -robust, because it has to contain \vec{s} or \vec{s} and so, by (P), cannot also contain both \vec{r} and $\vec{r} \wedge \vec{s}$ or both \vec{r} and $\vec{r} \wedge \vec{s}$. Clearly, n -robust k -profiles induce n -robust ℓ -profiles for all $\ell < k$, and every n -robust profile is also m -robust for every $m < n$.

Lemma 3.5. *Let n be a positive integer. Let $r \in U$ be a separation that efficiently distinguishes two n -robust profiles P, P' in \vec{U} , and let $s \in U$ be a separation that efficiently distinguishes two profiles \hat{P}, \hat{P}' in \vec{U} . If $|r| < |s| < n$, then r has an orientation \vec{r} such that either $\vec{r} \wedge \vec{s}$ or $\vec{r} \wedge \vec{s}$ efficiently distinguishes \hat{P} from \hat{P}' .*

Proof. As $|r| < |s|$ and s distinguishes \hat{P} from \hat{P}' , efficiently, \hat{P} and \hat{P}' also orient r , and they do so identically, say as \vec{r} . Let \vec{s} be the orientation of s in \hat{P} ; then $\vec{s} \in \hat{P}'$. If $|\vec{r} \wedge \vec{s}| \leq |s|$ then $\vec{r} \wedge \vec{s}$ distinguishes \hat{P} from \hat{P}' (and efficiently): it lies in \hat{P} by $\vec{r} \wedge \vec{s} \leq \vec{s} \in \hat{P}$ and the consistency of \hat{P} , but not in \hat{P}' by $\vec{r}, \vec{s} \in \hat{P}'$ and (P). Similarly, if $|\vec{r} \wedge \vec{s}| \leq |s|$ then $\vec{r} \wedge \vec{s}$ efficiently distinguishes \hat{P} from \hat{P}' . So let us assume that $|\vec{r} \wedge \vec{s}|, |\vec{r} \wedge \vec{s}| > |s|$, and derive a contradiction.

By submodularity, $\vec{r}_1 = \vec{r} \wedge \vec{s}$ and $\vec{r}_2 = \vec{r} \wedge \vec{s}$ have order $< |r|$. As r distinguishes P from P' , it has orientations $\vec{r} \in P$ and $\vec{r} \in P'$. As it does so efficiently, P and P' orient r_1 and r_2 identically. As P' contains $\vec{r}_1, \vec{r}_2 \leq \vec{r}$ by consistency, we thus have $\vec{r}_1, \vec{r}_2 \in P$ too, as well as $\vec{r} \in P$ by assumption. This contradicts the n -robustness of P , since $s \in S_n$. \square

The reader will have noticed that, in the proof of Lemma 3.5, we did not in fact use that P and P' themselves are n -robust: we only used the implied n -robustness of the k -profiles they induce for $k = |r| + 1$. Let us say that two profiles P, P' in \vec{U} are *n-robustly distinguishable* if they are distinguishable and the distinct k -profiles they induce for $k = \kappa(P, P') + 1$ are both n -robust.

A set \mathcal{P} of profiles in \vec{U} is *n-robust* if any distinct $P, P' \in \mathcal{P}$ are n -robustly distinguishable and satisfy $\kappa(P, P') < n$. If a set \mathcal{P} of distinguishable profiles is n -robust for

$$n = \max\{\kappa(P, P') \mid P, P' \in \mathcal{P}\} + 1,$$

we call it *robust*. This will be the case if every $P \in \mathcal{P}$ is robust and the profiles in \mathcal{P} are pairwise distinguishable. But there can be robust sets of profiles that are not individually robust, as their individual robustness may fail just for n much larger than the order of separations needed to distinguish \mathcal{P} .

Theorem 3.6. (Canonical tangle-tree theorem for separation universes)

*Let $\vec{U} = (\vec{U}, \leq, *, \vee, \wedge, | \cdot |)$ be a submodular universe of separations. Then for every robust set \mathcal{P} of profiles in \vec{U} there is a nested set $T = T(\mathcal{P}) \subseteq U$ of separations such that:*

- (i) every two profiles in \mathcal{P} are efficiently distinguished by some separation in T ;
- (ii) every separation in T efficiently distinguishes a pair of profiles in \mathcal{P} ;
- (iii) for every automorphism α of \vec{U} we have $T(\mathcal{P}^\alpha) = T(\mathcal{P})^\alpha$;
- (iv) if all the profiles in \mathcal{P} are regular, then T is a regular tree set.

Proof. For every \mathcal{P} and $k > 0$ let \mathcal{P}_k denote the set of all n -robust k -profiles induced by profiles in \mathcal{P} , where $n = \max\{\kappa(P, P') \mid P, P' \in \mathcal{P}\} + 1$. As \mathcal{P} is robust, distinct $P, P' \in \mathcal{P}$ induce distinct elements of \mathcal{P}_k for $k = \kappa(P, P') + 1$. Hence any set $T \subseteq U$ distinguishing every \mathcal{P}_k will also distinguish \mathcal{P} .

We shall construct, simultaneously for all \mathcal{P} , nested sets $T_0 \subseteq T_1 \subseteq \dots$ of separations $T_k = T_k(\mathcal{P})$ such that, for all $k \geq 1$:

- (I) T_k efficiently distinguishes \mathcal{P}_k ;
- (II) every $s \in T_k \setminus T_{k-1}$ efficiently distinguishes some pair of profiles in \mathcal{P}_k ;
- (III) for every automorphism α of \vec{U} we have $T_k(\mathcal{P}^\alpha) = T_k(\mathcal{P})^\alpha$.

Then, clearly, $T := \bigcup_{k=1,2,\dots} T_k$ will satisfy (i)–(iii) and be nested, as desired. Statement (iv) will be an immediate consequence of (ii), because small separations cannot distinguish regular profiles.

We start our construction with $T_0 = \emptyset$. Now consider $k > 0$, and assume that for all $0 < \ell < k$ and all \mathcal{P} we have constructed $T_\ell = T_\ell(\mathcal{P})$ so as to satisfy (I)–(III).

If $k = 1$, put $Q := \emptyset$ and $\mathcal{P}_Q := \mathcal{P}_1$, and let S_Q be the set of all \mathcal{P}_1 -relevant separations in S_1 . Clearly, S_Q distinguishes \mathcal{P}_Q efficiently.

If $k > 1$, suppose some distinct $P, P' \in \mathcal{P}_k$ are not yet distinguished by T_{k-1} . The $(k-1)$ -profiles that P and P' induce are again n -robust, so they lie in \mathcal{P}_{k-1} , and hence by (I) cannot be distinct. We denote the set of all profiles in \mathcal{P}_k that induce the same profile $Q \in \mathcal{P}_{k-1}$ by \mathcal{P}_Q . As \mathcal{P}_k is the union of all such \mathcal{P}_Q , it will suffice to find for every $Q \in \mathcal{P}_{k-1}$ a nested set $T_Q \subseteq U$ that efficiently distinguishes \mathcal{P}_Q , and make sure that these T_Q are also nested with T_{k-1} and with each other.

If $k > 1$, consider any fixed $Q \in \mathcal{P}_{k-1}$. Let S_Q be the set of all those separations $s \in S_k$ which distinguish some $\hat{P}, \hat{P}' \in \mathcal{P}_Q$ that are not distinguished by any separation in S_k that crosses fewer separations in T_{k-1} than s does. Clearly, every separation in S_Q is \mathcal{P}_Q -relevant, and S_Q distinguishes \mathcal{P}_Q efficiently.

Let us show that S_Q is nested with T_{k-1} . Suppose there is a separation $s \in S_Q$ that crosses a separation $r \in T_{k-1}$. Then $k > 1$, as $T_0 = \emptyset$. By the definition of S_Q , there exist $\hat{P}, \hat{P}' \in \mathcal{P}_Q$ that are distinguished by s but by no separation in S_k that crosses fewer separations in T_{k-1} than s does. By Lemma 3.5, r and s have orientations \vec{r} and \vec{s} such that $\vec{r} \wedge \vec{s}$ or $\vec{r} \wedge \vec{s}$ efficiently distinguishes \hat{P} from \hat{P}' . These corner separations of r and s are not only, unlike s , nested with r but also, by Lemma 2.1, with every separation $t \in T_{k-1}$ that is nested with s , because t is also nested with r since both lie in T_{k-1} . This contradiction to the choice of \hat{P} and \hat{P}' completes our proof that S_Q is nested with T_{k-1} .

Our next aim is to find a nested subset T_Q of S_Q that still distinguishes all of \mathcal{P}_Q . As S_Q distinguishes \mathcal{P}_Q , it suffices to find such a set T_Q that distinguishes $\mathcal{P}_Q \upharpoonright S_Q$. This will exist by Theorem 3.3 if S_Q scatters $\mathcal{P}_Q \upharpoonright S_Q$.

To prove this, consider two profiles $P, P' \in \mathcal{P}_Q \upharpoonright S_Q$ and two crossing separations $r, s \in S_Q$ with orientations $\vec{r}, \vec{s} \in P$ and $\vec{r}, \vec{s} \in P'$. Let us show that both $\vec{r} \wedge \vec{s}$ and $\vec{r} \vee \vec{s}$, like r and s , have order $k - 1$. Neither of them can have smaller order than $k - 1$: if one of them did then it would, by (P) and consistency, distinguish the profiles in \mathcal{P}_Q that induce P on S_Q from those that induce P' , so two such profiles in \mathcal{P}_Q would induce distinct profiles in \mathcal{P}_{k-1} , whereas in fact they all induce Q . Thus, $|\vec{r} \wedge \vec{s}|, |\vec{r} \vee \vec{s}| \geq k - 1$. Since $|\vec{r} \wedge \vec{s}| + |\vec{r} \vee \vec{s}| \leq |r| + |s|$ by submodularity, and $|r| = |s| = k - 1$ as $r, s \in S_Q$, we thus have $|\vec{r} \wedge \vec{s}| = |\vec{r} \vee \vec{s}| = k - 1$ as claimed.

Thus, $\vec{r} \vee \vec{s}$ lies in \vec{S}_k , and it distinguishes P from P' . By Lemma 2.1 it is nested with every $t \in T_{k-1}$, because $r, s \in S_Q$ are nested with T_{k-1} (as shown earlier). By definition of S_Q , this means that $\vec{r} \vee \vec{s}$ lies in \vec{S}_Q . But then it also lies in P , by $\vec{r}, \vec{s} \in P$ and (P). This completes our proof that S_Q scatters \mathcal{P}_Q .

Applying Theorem 3.3 to $\mathcal{P}_Q \upharpoonright S_Q$ in $(\vec{S}_Q, \leq, *)$ for each $Q \in \mathcal{P}_{k-1}$ (respectively, for $Q = \emptyset$ if $k = 1$), we obtain a family of sets $T_Q \subseteq S_Q$ of separations distinguishing \mathcal{P}_Q . Each of these T_Q is nested, and nested with T_{k-1} . In order to show that they are also nested with each other, consider any $s \in T_Q$ and $s' \in T_{Q'}$ for distinct $Q, Q' \in \mathcal{P}_{k-1}$. By (I) for $k - 1$ there is a separation $r \in T_{k-1}$ with orientations $\vec{r} \in Q$ and $\vec{r} \in Q'$. Since T_Q is nested with T_{k-1} , the separations r and s have comparable orientations. So either \vec{r} or \vec{r} points towards s ; let us show that \vec{r} does.

Since $s \in S_Q$ it is \mathcal{P}_Q -relevant, so it has orientations $\vec{s} \in P_1$ and $\vec{s} \in P_2$ for some $P_1, P_2 \in \mathcal{P}_Q$. As P_1, P_2 induce Q , they both contain \vec{r} . Hence if \vec{r} points towards s , then either $\vec{r} < \vec{s}$ or $\vec{r} < \vec{s}$, contradicting the consistency of P_1 or P_2 , respectively. So \vec{r} points towards s , as claimed.

Similarly, \vec{r} points towards s' . We thus have $\vec{s}' < \vec{r} < \vec{s}$ for suitable orientations of s and s' . In particular, s and s' are nested.

We have shown that the T_Q are nested sets of separations that are also nested with each other and with T_{k-1} . Hence, $T_k := T_Q$ for $k = 1$, and otherwise

$$T_k := T_{k-1} \cup \bigcup_{Q \in \mathcal{P}_{k-1}} T_Q$$

is nested too, and by construction it satisfies (I) and (II).

In order to verify (III), let α be an arbitrary automorphism of $(\vec{U}, \leq, *, \vee, \wedge, | |)$. If $k = 1$, then (III) holds by definition of T_1 , which we obtained from Theorem 3.3 as a canonical tree set $T \subseteq S_1$ distinguishing \mathcal{P}_1 . Now consider $k > 1$. Note that $\mathcal{P}_\ell^\alpha = (\mathcal{P}_\ell)^\alpha$ for all ℓ , by definition of \mathcal{P}_ℓ . Hence $S_{Q^\alpha} = (S_Q)^\alpha$ for every $Q \in \mathcal{P}_{k-1}$, by definition of S_Q for \mathcal{P} and S_{Q^α} for \mathcal{P}^α , and (III) for $k - 1$. Therefore $\mathcal{P}_{Q^\alpha}^\alpha \upharpoonright S_{Q^\alpha} = (\mathcal{P}_Q \upharpoonright S_Q)^\alpha$, so Theorem 3.3 implies that $T(\mathcal{P}_{Q^\alpha}^\alpha \upharpoonright S_{Q^\alpha}) =$

$T(\mathcal{P}_Q \upharpoonright S_Q)^\alpha$, for all $Q \in \mathcal{P}_{k-1}$. Hence, again using (III) for $k-1$,

$$\begin{aligned} T_k(\mathcal{P}^\alpha) &= T_{k-1}(\mathcal{P}^\alpha) \cup \bigcup_{Q^\alpha \in \mathcal{P}_{k-1}^\alpha} T(\mathcal{P}_{Q^\alpha}^\alpha \upharpoonright S_{Q^\alpha}) \\ &= T_{k-1}(\mathcal{P})^\alpha \cup \bigcup_{Q \in \mathcal{P}_{k-1}} T(\mathcal{P}_Q \upharpoonright S_Q)^\alpha \\ &= T_k(\mathcal{P})^\alpha, \end{aligned}$$

completing the proof of (I)–(III) for k . \square

Let us add a few remarks. First, we stated (iii) for automorphisms rather than arbitrary isomorphism of universes just to keep things simple; it clearly generalizes as in Theorem 3.3 (whose more general form we needed in the proof of Theorem 3.6).

Since trivial separations cannot distinguish profiles, one might expect that, by (ii), T will always be a tree set. However this is not the case. For example, \vec{T} might contain a small separation \vec{r} that distinguishes two k -profiles in \mathcal{P} . This \vec{r} will not be trivial in S_k . But it may be trivial in some S_n with $n > k$. Then no n -profile in \vec{U} will contain \vec{r} , but r may still be needed for T_k . Under mild additional assumptions,⁸ however, one can replace the premise in (iv) with the weaker premise that no $P \in \mathcal{P}$ is special.

In applications, the robustness of \mathcal{P} will often be a consequence of the fact that the profiles it contains are themselves robust, and often they will also be regular. Here is a leaner version of Theorem 3.6 for robust regular profiles:

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

- (i) *every two profiles in \mathcal{P} are efficiently distinguished by some separation in T ;*
- (ii) *every separation in T efficiently distinguishes a pair of profiles in \mathcal{P} ;*
- (iii) *for every automorphism α of \vec{U} we have $T(\mathcal{P}^\alpha) = T(\mathcal{P})^\alpha$.* \square

But note that we cannot simply omit all robustness requirements in Theorem 3.6 and Corollary 3.7; see [4] for a counterexample.

As a common basis for the proof of our Theorems 3 and 4, let us note a further simplification of Theorem 3.6. Rather than starting with any fixed set \mathcal{P} of distinguishable profiles, we find a tree set that distinguishes them all:

Corollary 3.8. *Let $\vec{U} = (\vec{U}, \leq, *, \vee, \wedge, | \cdot |)$ be a submodular universe of separations. There is a canonical regular tree set $T \subseteq \vec{U}$ that efficiently distinguishes all the distinguishable robust and regular profiles in \vec{U} .*

⁸For example, it suffices to assume that, when \vec{s} is small, $\vec{r} \leq \vec{s}$ implies $|r| \leq |s|$. This holds for separations of graphs and matroids.

Proof. Let \mathcal{P}' be the set of robust and regular profiles in \vec{U} , and let \mathcal{P} be the set of its inclusion-maximal elements. Note that the profiles in \mathcal{P} are pairwise distinguishable. By Corollary 3.7, there is a canonical tree set $T \subseteq \vec{U}$ that efficiently distinguishes \mathcal{P} . We have to show that it efficiently distinguishes every pair P_1, P_2 of distinguishable profiles in \mathcal{P}' .

Each P_i is a k_i -profile for some k_i . By definition of \mathcal{P} it extends to some $P'_i \in \mathcal{P}$, which is a k'_i -profile for some $k'_i \geq k_i$. (Recall that all the profiles we are considering are profiles ‘in’ \vec{U} , and hence k -profiles for some k .) Since T distinguishes the P'_i efficiently, and any separation distinguishing the P_i will also distinguish their extensions P'_i , our assumption that the P_i are distinguishable means that T distinguishes the P'_i by a separation of order at most $\kappa(P_1, P_2)$. This separation, therefore, has orientations in both P_1 and P_2 , and thus distinguishes these efficiently too. \square

3.3 Minimal distinguishing tree sets, non-canonical

As shown by Example 4 we cannot, without losing canonicity, strengthen (ii) in Theorem 3.6 to saying that T is minimal given (i). But if we are prepared to make do without canonicity, we can strengthen (ii) even more: we can ensure that T is minimal not only with (i) but more generally with the property that it distinguishes \mathcal{P} , efficiently or not. (In fact, we shall prove that every T that is minimal with (i) is minimal even as a \mathcal{P} -distinguishing set of separations; see Lemma 3.10 below.)

Robertson and Seymour [20] proved that every graph has a (non-canonical) tree-decomposition that distinguishes any given set Θ of distinguishable tangles, in the sense that any distinct $\theta, \theta' \in \Theta$ disagree on some separation associated with the tree-decomposition. Moreover, this decomposition has only $|\Theta|$ parts. It follows that every tangle $\theta \in \Theta$ ‘lives in’ one of these parts – the unique part towards which θ orients all the separations associated with the tree-decomposition that it orients – and every part is home to a tangle from Θ .

Although in our abstract separation system we have no vertices, and our tree set T has no ‘parts’, we can generalize this to a non-canonical version of our theorem by representing those parts as consistent orientations of T . For every $P \in \mathcal{P}$, the subset $P \cap \vec{T}$ of \vec{T} is consistent, and hence extends to a consistent orientation of all of T .⁹ This extension is not in general unique, but for our minimal T it will be. Moreover, every consistent orientation of T can be obtained from some $P \in \mathcal{P}$ in this way. In this sense, every profile in \mathcal{P} will ‘live in’ a unique consistent orientation of T , and every consistent orientation of T will be home to some $P \in \mathcal{P}$.

Theorem 3.9. (Non-canonical tangle-tree theorem for separation universes)
*Let $\vec{U} = (\vec{U}, \leq, *, \vee, \wedge, | \cdot |)$ be a submodular universe of separations. For every robust set \mathcal{P} of regular profiles in \vec{U} there is a regular tree set $T' \subseteq \vec{U}$ of separations such that:*

⁹We need here that no element of $P \cap \vec{T}$ is co-trivial in T , which we shall ensure by requiring \mathcal{P} to be regular. See [6, Lemma 4.1] for a formal proof that such extensions exist.

- (i) every two profiles in \mathcal{P} are efficiently distinguished by some separation in T' ;
- (ii) no proper subset of T' distinguishes \mathcal{P} ;
- (iii) for every $P \in \mathcal{P}$ the set $P \cap \vec{T}'$ extends to a unique consistent orientation $P_{T'}$ of T' . This map $P \mapsto P_{T'}$ from \mathcal{P} to the set \mathcal{O} of consistent orientations of T' is bijective.

Note that, by (i), assertion (ii) implies that every separation in T' distinguishes two profiles in \mathcal{P} efficiently, as earlier in Theorem 3.6.

Also, the implicit assumption in the theorem that the profiles in \mathcal{P} must be pairwise distinguishable (because we require \mathcal{P} to be robust) is not a real restriction. Indeed, the tree set T' returned by Theorem 3.9 will distinguish *all* the distinguishable robust and regular profiles in \vec{U} : just take as \mathcal{P} the set of all the maximal such profiles, and argue as in the proof of Corollary 3.8.

For the proof of Theorem 3.9 we need a fun little lemma about separating edge sets in trees, which may also be of interest in its own right. The intended application is that choosing T' in Theorem 3.6 minimal with respect to (i), ignoring canonicity, will automatically make it satisfy (ii) of Theorem 3.9. Recall that a *U-path* in a graph G with $U \subseteq V(G)$ is a path that meets U in exactly its distinct ends.

Lemma 3.10. *Let G be a tree with an edge labelling $\ell: E(G) \rightarrow \mathbb{N}$. Let $U \subseteq V(G)$ and $F \subseteq E(G)$ be such that every U -path $P \subseteq G$ contains an edge $e \in F$ such that $\ell(e) = \min\{\ell(e') \mid e' \in P\}$ and F is minimal with this property, given U . Then for every $e \in F$ there is a U -path in G whose only edge in F is e .*

Proof. Let us call an edge $e \in G$ *essential* on a path $P \subseteq G$ if e is the unique edge of P in F with $\ell(e) = \min\{\ell(e') \mid e' \in P\}$. The minimality of F assumed in the lemma means that every edge in F is essential on some U -path. If the lemma fails, there is an edge $e \in F$ that is not the only F -edge on any U -path. Let such an edge e be chosen with $\ell(e)$ maximum, and let $P = uGv$ be a U -path on which e is essential.




Figure 2: Paths in the proof of Lemma 3.10

By definition, e is not the only edge of P in F ; choose another, f say, as close to e as possible. Assume that f lies on the segment ePv of P strictly between e and v . Then $\ell(f) > \ell(e)$ since e is essential on P , so by the choice of e there exists a U -path $Q = xGy$ whose only edge in F is f . Let x be the end of Q for which $f \notin eGx$. By the choice of f , this path $eGx \subseteq P \cup Q$ has no edge in F .

If uPe has no edge in F , then e is the unique F -edge on the U -path uGx , contradicting the choice of e . So uPe contains an edge $f' \in F$. Choose f' as close to e as possible, and repeat the earlier argument with f' in the place of f to find a U -path $Q' = x'Gy'$ whose only edge in F is f' , with $f' \notin eGx'$ say. By the choice of f' , this path $eGx' \subseteq P \cup Q'$ has no edge in F . Hence e is the only F -edge on the U -path xGx' , contradicting the choice of e . \square

Proof of Theorem 3.9. Let T be the regular tree set provided for \mathcal{P} by Theorem 3.6. Then $T' = T$ satisfies (i) of our theorem. Now choose $T' \subseteq T$ minimal with this property.

For every $P \in \mathcal{P}$, the set $P \cap \vec{T}'$ extends uniquely to a consistent orientation $P_{T'}$ of T' . (4)

To prove this, note first that $P \cap \vec{T}'$ is consistent, because P is. By [6, Lemma 4.1] and our assumption that all $P \in \mathcal{P}$ are regular, $P \cap \vec{T}'$ extends to a consistent orientation $P_{T'}$ of all of T' .

For a proof that $P_{T'}$ is unique, let $\ell := \max\{|\vec{r}| : \vec{r} \in P \cap \vec{T}'\}$. Since P is an m -profile, for some $m > \ell$, it orients *all* the separations of order $\leq \ell$ in T . If $P \cap \vec{T}'$ extends to distinct consistent orientations of T' there is some $s \in T'$ which P does not orient, which thus has order $k := |s| > \ell$. In fact, s is not oriented by the down-closure $\lceil P \cap \vec{T}' \rceil$ of $P \cap \vec{T}'$ in \vec{T}' , because this is contained in all consistent orientations of T' extending $P \cap \vec{T}'$, by definition of consistency.

When s was added to T in its construction, the reason was that s distinguishes some profiles $P', P'' \in \mathcal{P}$ inducing the same profile Q of S_k , with $\vec{s} \in P'$ and $\vec{s} \in P''$ say. Notice that every $\vec{r} \in P \cap \vec{T}'$ points to s : as T' is nested, \vec{r} is either smaller or greater than some orientation of s , but s has no orientation in $\lceil P \cap \vec{T}' \rceil \supseteq \lceil \vec{r} \rceil$.

Hence for every $\vec{r} \in P \cap \vec{T}'$, which lies in \vec{T}_k by definition of ℓ and k , we have either $\vec{r} < \vec{s}$ and thus $\vec{r} \in P' \cap \vec{T}_k \cap \vec{T}' = Q \cap \vec{T}'$ by the consistency of P' and $\vec{s} \in P'$, or else $\vec{r} < \vec{s}$ and $\vec{r} \in P'' \cap \vec{T}_k \cap \vec{T}' = Q \cap \vec{T}'$ by the consistency of P'' and $\vec{s} \in P''$. So on the separations in T' that it orients, P agrees with Q , and hence with both P' and P'' . Thus, T' does not distinguish P from either P' or P'' . But T' does distinguish all distinct profiles in \mathcal{P} . So P is not distinct from either P' or P'' , i.e., $P' = P = P''$ contradicting the choice of P' and P'' .

This completes our proof of (4). Thus, $P \mapsto P_{T'}$ is a well-defined map from \mathcal{P} to the set \mathcal{O} of consistent orientations of T' . Let us use this to prove assertion (ii) of our theorem.

As shown in [7], there is a tree G with edge set T' whose natural ordering on \vec{T}' , as defined in Section 2.3, coincides with our given ordering on \vec{T}' . As with every finite tree, the consistent orientations of its edge set with respect to this natural ordering correspond to its nodes: an orientation of $E(G)$ is consistent if and only if it orients all of $E(G)$ towards some fixed node of G [3, 7]. Since the consistent orientations of $E(G)$ are those of T' , we may thus think of \mathcal{O} as the node set of G .

Let us apply Lemma 3.10 to $G = (\mathcal{O}, T')$ with $F = T'$ and $U = \{P_{T'} \mid P \in \mathcal{P}\}$ and $\ell(s) := |s|$ for all $s \in T'$. Note that a separation $s \in T'$ distinguishes two

profiles $P, P' \in \mathcal{P}$ if and only if s lies on the unique $P_{T'}-P'_{T'}$ path in G : it is then, and only then, that orienting it towards the node $P_{T'}$ differs from orienting it towards the node $P'_{T'}$, and as mentioned earlier, an edge $\vec{s} \in \vec{T}$ is oriented towards a node $P_{T'}$ of G if and only if \vec{s} lies in $P_{T'}$.

Let us show that $F = T'$ satisfies the premise of the lemma. By definition of T' , every two profiles $P, P' \in \mathcal{P}$ are distinguished by some $s \in T'$ with $\ell(s) = \kappa(P, P')$, but as soon as we delete an element s of T' it loses this property. Then there are $P, P' \in \mathcal{P}$ such that s lies on the $P_{T'}-P'_{T'}$ path in G but every other edge r on that path has a label $\ell(r) = |r| > \kappa(P, P') = |s| = \ell(s)$. After deleting s from F no edge of minimum label on that path will lie in F . Hence, F is as required in the lemma.

The lemma asserts that, therefore, each $s \in T'$ is the only edge in T' that distinguishes some two profiles in \mathcal{P} depending on s . This is statement (ii) of our theorem. Moreover, it implies that $U = V(G)$: if s is the only edge in $F = T' = E(G)$ on some U -path, then this path has length 1, and hence the ends of s lie in U . Thus, our map $P \rightarrow P_{T'}$ from \mathcal{P} to \mathcal{O} is surjective. Since it is also injective, by (i), our proof is complete. \square

3.4 Symmetric profiles

This section is for those readers that have become interested in abstract profiles and wonder whether the concept might still be improved. Readers primarily interested in their applications – to graphs, matroids and elsewhere – may skip ahead to Section 4 without loss.

Our motivation for introducing the notion of a profile was to find a concept, as general as possible, that captures the idea of identifying highly connected parts in a discrete structure by orienting its low-order separations consistently towards it. While initially these orientations might have been induced by some concrete highly connected structure such as a k -block, and receive their consistency from there, the idea was that they might ultimately be thought of as a highly connected substructure in their own right, thus allowing us to treat these in more general contexts such as abstract separation systems.

The key to this, then, lies in finding an abstract notion of consistency to make this work. Our formal definition of consistency, that two separations should not point away from each other, is just a minimum requirement one would naturally make. Conditions (P) and (R) go a little further but are still fairly general; recall that tangles satisfy both. As shown in [4], there is little hope of weakening (R) and retaining a tangle-tree theorem. But maybe we can weaken (P) a little, while still keeping it narrow enough that the ‘weak profiles’ it then defines can still be distinguished in a tree-like way?

There is also the question of just how natural is the notion of a profile. Condition (P) seems very natural as a requirement of consistency: if \vec{r} and \vec{s} both point to some substructure X then so should $\vec{t} := \vec{r} \vee \vec{s}$ if t is oriented at all, provided the universe in which $\vec{r} \vee \vec{s}$ is taken is dense enough that there is ‘no room for X ’ between \vec{r} and \vec{s} on the one hand and $\vec{r} \wedge \vec{s} = \vec{t}$ on the other (cf. Example 1). And surely, then, every $\vec{t} \leq \vec{r} \vee \vec{s}$ should also lie in

such a profile P , as soon as t – rather than $\vec{r} \vee \vec{s}$ – is oriented at all, even if $\vec{r} \vee \vec{s}$ is not. Our current definition ensures this (by consistency) only if $\vec{r} \vee \vec{s}$ itself is oriented too, and hence lies in P .

Let us call a consistent orientation P of a separation system $(\vec{S}, \leq, *)$ in some universe $\vec{U} = (\vec{U}, \leq, *, \vee, \wedge)$ a *strong profile* if it satisfies this strengthening of (P):

$$\text{For all } \vec{r}, \vec{s} \in P \text{ and } \vec{t} \leq \vec{r} \vee \vec{s} \text{ the separation } \vec{t} \text{ is not in } P. \quad (\text{P}^+)$$

Surprisingly, perhaps, if \vec{U} is a distributive as a lattice and \vec{S} is *submodular* in the sense that, whenever $\vec{r}, \vec{s} \in \vec{S}$, either $\vec{r} \vee \vec{s} \in \vec{S}$ or $\vec{r} \wedge \vec{s} \in \vec{S}$ [6], then all its profiles are strong:

Theorem 3.11. [12, Theorem 1] *If \vec{U} is distributive and \vec{S} is submodular, then an orientation of S satisfies (P⁺) if and only if it is consistent and satisfies (P).*

We shall apply Theorem 3.11 in Section 4.2 to show that, with one simple exception, all profiles in a matroid are in fact tangles.

So strengthening profiles in this way yields little new – and anyway, our aim was to weaken (P), not strengthen it. But we can use the same idea to do that too, and in a way that also will make the notion of a profile more symmetric.

Condition (P⁺) says that P contains no triple $(\vec{r}_1, \vec{r}_2, \vec{r}_3)$ such that $\vec{r}_3 \leq \vec{r}_1 \vee \vec{r}_2$. To arrive at a weaker notion, let us ban a triple $(\vec{r}_1, \vec{r}_2, \vec{r}_3)$ only if it satisfies this for every permutation. More precisely, let us say that $\vec{r}_1, \vec{r}_2, \vec{r}_3 \in \vec{U}$ form a *bad triple* if

$$\vec{r}_k \leq \vec{r}_i \vee \vec{r}_j \quad \text{whenever} \quad \{i, j, k\} = \{1, 2, 3\}$$

(Figure 3), and call a consistent orientation of $S \subseteq U$ a *weak S-profile* if it contains no bad triple.¹⁰




Figure 3: A bad triple of separations

However, Eberenz [12] showed that weak profiles can no longer be separated in a tree-like way, as is our aim in this paper. Indeed, we have the following example:

¹⁰One can, of course, construct separation systems with profiles that are not weak profiles in this sense. However, there are no natural ones: by Theorem 3.11, profiles of submodular separation systems in distributive universes – which covers all natural examples – are also weak profiles.

Example 6. Let S consists of two crossing separations r, s and their four corner separations. Let all these six separations have the same order. Let $O \subseteq \vec{S}$ be the set the four corner separations oriented ‘towards the centre’, that is, the set of separations of the form $\vec{r} \wedge \vec{s}$, one for every choice of orientations \vec{r} of r and \vec{s} of s . Each of the four possible orientations \vec{r} and \vec{s} of r and s extends O to an orientation $P = P_{\vec{r}, \vec{s}}$ of S . If \vec{S} , together with a greatest element 1 and a least element 0, is the entire universe in which \vee and \wedge are defined, then each of these P will contain a bad triple of the form $(\vec{r}, \vec{s}, \vec{r} \wedge \vec{s})$, so none of them is a weak profile.

However, let us choose a graph G so that S becomes its set of proper 3-separations, and let \vee and \wedge be defined as usual in the universe of all separations of G . Then the four orientations P of S defined above are weak 3-profiles. To see this, notice that if \vec{x} is a corner separation and \vec{y} is any of the other separations in P , then $\vec{x} \vee \vec{y}$ has order at least 4 unless $\vec{x} \leq \vec{y}$, and r has no orientation $\vec{r} \leq \vec{x} \vee \vec{y}$ (and likewise for s). Here, $\vec{x} \vee \vec{y} \notin \vec{S}$, but allowing this, while still banning $\vec{r} \in \vec{S}$ from P if $\vec{r} \leq \vec{x} \vee \vec{y}$ and $\vec{x}, \vec{y} \in P$, was the whole point of introducing bad triples and weak profiles.

Now we cannot distinguish these four weak profiles by a nested set of separations. Indeed, r and s are the only separations that can distinguish any of them, we cannot have both r and s in a nested set, but one of them will not distinguish the four weak profiles pairwise.

4 Applications to graphs and matroids

In this section we apply the results of Section 3 to graphs and matroids, concluding with the proofs of the theorems stated in the introduction. We use the terms ‘block’, ‘profile’ and ‘tangle’ for k -blocks, k -profiles and k -tangles with any k . By a profile *in* a graph or matroid we mean a profile in the universe of all its separations.

4.1 Canonical tangle-tree theorems in graphs

Let $G = (V, E)$ be a finite graph, and let $\vec{U} = (\vec{U}, \leq, *, \vee, \wedge, | |)$ be the universe of all its oriented separations (A, B) with the usual submodular order function of $|(A, B)| := |A, B| := |A \cap B|$. We think of $\{A, B\}$ as the corresponding unoriented separation, which officially in our terminology is the set $\{(A, B), (B, A)\}$.

Given a k -block b in G , let us write

$$P_k(b) := \{(A, B) \in \vec{S}_k \mid b \subseteq B\}$$

for the orientation of S_k it induces. In the introduction, we called this the ‘profile of b ’. Our later definition of abstract profiles bears this out:

Lemma 4.1. *$P_k(b)$ is a regular k -profile.* □

We shall call the profiles of the form $P_k(b)$ the *(k -)block profiles* in G .

Notice that a separation $(A, B) \in S_k$ distinguishes two k -blocks, in the original sense that one lies in A and the other in B , if and only if (A, B) distinguishes the k -profiles they induce. From now on, we shall use the term ‘distinguish’ formally only for profiles, and say that a separation *distinguishes* two blocks, say, if it distinguishes their profiles. In this way, we can now also speak of separations distinguishing blocks from other profiles such as tangles.

Similarly, we shall call a block b *robust* if the profiles¹¹ it induces are robust, and similarly for sets of blocks.

A *tangle of order k* in G , or *k -tangle*, is an orientation θ of S_k such that

$$G[A_1] \cup G[A_2] \cup G[A_3] \neq G \quad (\text{T})$$

for all triples $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \theta$. These (A_i, B_i) need not be distinct, so tangles are consistent. In fact, the following is immediate from the definitions:

Lemma 4.2. *Every tangle in G is a robust regular profile.* □

In order to apply Theorem 3.6 to prove Theorem 3, we need to show that every tree set of separations in a graph G is induced by a suitable tree-decomposition of G . The following lemma, which was proved for regular tree sets in [4] and for arbitrary tree sets in [7],¹² achieves this:

Lemma 4.3. *Let T be a tree set of separations of G . Then G has a tree-decomposition $(\mathcal{T}, \mathcal{V})$ such that:*

- (i) *T is precisely the set of separations of G associated with the edges of \mathcal{T} ;*
- (ii) *if T is invariant under the automorphisms of G , then $(\mathcal{T}, \mathcal{V})$ is canonical.*¹³

Note that we have an immediate proof of Theorem 2 now. Indeed, since k -profiles are n -robust for all $n \leq k$, the set \mathcal{P} of all k -block profiles and k -tangles in a graph G , for k fixed, is always robust. Theorem 3.6 therefore yields a regular and invariant tree set T distinguishing all its k -blocks and k -tangles, which by Lemma 4.3 yields the desired tree-decomposition of G .

Theorem 3 follows in a similar way from Corollary 3.8. This takes care of the fact that, when k is variable, profiles can induce each other:

Theorem 3. *Every finite graph G has a canonical tree-decomposition that efficiently distinguishes all its distinguishable tangles and robust blocks.*

Proof. By Lemmas 4.1 and 4.2, all the profiles in G that are tangles or induced by a robust block are robust and regular profiles. By Corollary 3.8 there is a canonical regular tree set T that efficiently distinguishes every two of them. Lemma 4.3 turns T into a tree-decomposition of G . □

¹¹Note that b can be a k -block for several k . If the corresponding sets S_k are distinct, then so are the k -profiles $P_k(b)$ that b induces, though clearly $P_\ell(b) \subseteq P_k(b)$ for $\ell < k$.

¹²Part (ii) of Lemma 4.3 was not stated in the sources cited, but the proofs given establish it.

¹³Canonical tree-decompositions of graphs are defined in the Introduction.

Let us conclude this section with two observations about profiles in graphs. The first is a kind of ‘inverse consistency’ phenomenon. If P is a k -profile in a graph, $(A, B) \in P$, and $\{A', B'\} \in S_k$ is such that $A' \subseteq A$ and $B' \supseteq B$, then also $(A', B') \in P$, by the consistency of P and $(A', B') \leq (A, B)$.

But now suppose that, instead, $A' \supseteq A$ and $B' \subseteq B$. Then $(A, B) \leq (A', B')$, but still we get $(A', B') \in P$ if P is regular¹⁴ and $\{A', B'\}$ does not differ from $\{A, B\}$ too much:

Proposition 4.4. *Let P be a regular k -profile in a graph G , and let $(A, B) \in P$. Let $A' \supseteq A$ and $B' \subseteq B$ be such that both $\{A, B'\}$ and $\{A', B\}$ lie in S_k . Then $(A', B') \in P$.*

Proof. Let us first prove the assertion for $A' = A$, i.e., show that $(A, B') \in P$. We show that if this fails then $\{(A, B), (B', A), (B \cap A, A \cup B')\} \subseteq P$, in violation of (P).

We have $\{A, B'\} \in S_k$ by assumption, so if (A, B') does not lie in P as claimed, then (B', A) does. And since $A \cup B' = V$, we have $\{B \cap A, A \cup B'\} \in S_k$. Since P is regular, it must contain the small separation $(B \cap A, A \cup B')$.

The case of $A' \supseteq A$ can now be derived from this as follows. If $(A', B') \notin P$ then $(B', A') \in P$, since $\{A', B'\} \in S_k$. Then $(B', A) \in P$ by the above case applied to (B', A') , since $A \subseteq A'$. But this contradicts the fact that, as established above, $(A, B') \in P$. \square

Our second observation is that graphs can have profiles that are neither block profiles nor tangles:

Example 7. (Carmesin, personal communication 2014) The graph in Figure 4 has a 5-profile that orients both its 4-separations towards the middle. This profile is neither a 5-tangle nor induced by a 5-block.




Figure 4: A 5-profile that is neither a 5-tangle nor induced by a 5-block.

Example 7 makes it desirable to have a tangle-tree theorem for all the profiles in a graph, not just those that are tangles or block profiles. Corollary 3.8 yields this too, even without a need to require regularity.

¹⁴Regularity is not a severe restriction. We shall see in a moment that k -profiles in graphs are all regular for $k \geq 3$, and for $k = 1, 2$ the few irregular k -profiles are exactly understood.

To see this, we need some observations from [8]. First, that all k -profiles in a graph $G = (V, E)$ are regular as soon as $k > 2$. For $k = 1$ there can be exactly one irregular 1-profile, the set $\{(V, \emptyset)\}$, and only if G is connected. For $k = 2$ there can be irregular profiles, but these can be described precisely: they are precisely the 2-profiles the form

$$P_v = \{(A, B) \in \vec{S}_2 \mid v \in B \text{ and } (A, B) \neq (\{v\}, V)\}$$

where $v \in V$ is not a cutvertex of G . Let us call such profiles in graphs *principal*. Thus, all profiles in graphs other than those 2-profiles P_v are non-principal.

Theorem 4.5. *Every finite graph has a canonical tree-decomposition that distinguishes all its distinguishable robust non-principal profiles.*

Proof. By Corollary 3.8 there is a canonical regular tree set T that distinguishes all the distinguishable, robust, non-principal k -profiles in G for $k > 1$ as well as for $k = 1$ if they are regular, since all these profiles for $k > 1$ are regular [8]. Lemma 4.3 turns T into a tree-decomposition $(\mathcal{T}, \mathcal{V})$ that distinguishes all these profiles, with $\mathcal{V} = (V_t)_{t \in \mathcal{T}}$ say. If G is disconnected, then these are all its robust non-principal profiles, and the proof is complete.

If G is connected, it also has the irregular non-principal 1-profile $\{(V, \emptyset)\}$. To accommodate this, we add a new empty part V_t to our decomposition, joining t to some node z of \mathcal{T} by an edge associated with the separation $\{\emptyset, V\}$. This separation will distinguish $\{(V, \emptyset)\}$ from all the other profiles in G , since these contain (\emptyset, V) by consistency.

The only problem is how to choose z so that the tree-decomposition remains canonical. If \mathcal{T} has a central node, we choose this as z . If not, it has a central edge xy . Subdivide xy by a new node z , with $V_z := V_x \cap V_y$. Then every automorphism of \mathcal{T} is, or extends to, an automorphism of this (possibly modified) tree \mathcal{T}' that fixes z . In particular, the automorphisms of G , which act on \mathcal{T} by Lemma 4.3, still act on \mathcal{T}' . Now join t to z . Since all automorphisms of G fix $V_t = \emptyset$, the extended tree-decomposition will again be canonical. \square

4.2 The canonical tangle-tree theorem for matroids

Let M be a finite matroid with ground set E . Let $\vec{U} = (\vec{U}, \leq, *, \vee, \wedge, | |)$ be the universe of the oriented bipartitions of E , where \vee and \wedge are defined as for set separations in Section 2.2, and $| |$ is the usual submodular order function given by

$$|\{A, B\}| = |(A, B)| = r(A) + r(B) - r(E) + 1,$$

where r is the rank function of M . As before, we write \vec{S}_k for the separation system of all the separations in \vec{U} of order less than k .

Tangles in matroids were introduced implicitly by Robertson and Seymour [20], and explicitly by Geelen, Gerards, Robertson and Whittle [14]. A k -tangle in M is an orientation of S_k that has no element of the form $(E \setminus \{e\}, e)$ and no subset of the form

$$\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\} \text{ with } A_1 \cup A_2 \cup A_3 = E. \quad (\text{T}_M)$$

As with graphs, the separations $(A_1, B_1), (A_2, B_2), (A_3, B_3)$ above need not be distinct. In particular, tangles are consistent. Since they obviously also satisfy (P), they are in fact profiles, indeed robust and regular profiles.

Unlike for graphs, however, the tangles in a matroid are essentially all its profiles. Let us call a profile P in M *principal* if there exists an $e \in E$ such that $(A, B) \in P$ if and only if $e \in B$, for all bipartitions $\{A, B\}$ oriented by P .

Lemma 4.6. *For every integer $k \geq 1$, the k -tangles in a matroid are precisely its non-principal k -profiles.*

Proof. We have seen that matroid tangles are profiles, and they are non-principal by definition. Conversely note that, by consistency, a profile is principal as soon as it contains a separation of the form $(E \setminus \{e\}, \{e\})$. Hence a non-principal profile contains no such separation, and thus satisfies the tangle axiom requiring this. And by Theorem 3.11 it satisfies (P^+) , which immediately implies (T_M) . \square

A *tree-decomposition* of M is a pair (T, \mathcal{V}) where T is a tree and $\mathcal{V} = (V_t \mid t \in V(T))$ is a partition of E . As in graphs, we say $(\mathcal{T}, \mathcal{V})$ (*efficiently*) *distinguishes* two tangles in M if the set of separations of M that are associated with the edges of T – defined as for graphs – distinguishes these tangles (efficiently). Lemmas 4.2 and 4.3 hold for tree-decompositions of matroids too, with the same easy proofs.

Theorem 4 follows from Corollary 3.8 much as Theorem 3 did:

Theorem 4. *Every finite matroid has a canonical tree-decomposition which efficiently distinguishes all its distinguishable tangles.*

Proof. By Lemma 4.2, the tangles in M are robust and regular profiles. By Corollary 3.8 there is a canonical regular tree set T that efficiently distinguishes every two of them. Lemma 4.3 turns T into a tree-decomposition of M . \square

5 Applications in data analysis and optimization

5.1 Applying profiles to cluster analysis

Both the canonical and the non-canonical tangle-tree theorem lend themselves to the analysis of big data sets when clusters in these are interpreted as profiles. This approach to clusters has an important advantage over more traditional ways of identifying clusters: real-world clusters tend to be fuzzy, and profiles can capture them despite their fuzziness; cf. the grid example from the introduction.

Let D be a (large data) set with at least two elements, and let U be a set of separations $\{A, B\}$ of D whose orientations (A, B) form a universe \vec{U} of set separations as defined in Section 2.1.¹⁵ Let us call D *submodular* if it comes with a fixed submodular such universe \vec{U} , one with a submodular order function $\{A, B\} \mapsto |A, B|$ on U .

¹⁵We allow A and B to be empty, since $\{\emptyset, D\} \in U$ is needed for \vec{U} to be a universe.

How exactly this order function should be defined will depend both on the type of data that D represents and on the envisaged type of clustering. See [11] for examples of how this might be done when D is the set of pixels of an image, U consists of just the bipartitions of D , and the clusters to be captured are the natural regions of this image such as a nose, or a cheek, in a portrait.

If U consists of just the bipartitions of D , then all profiles in \vec{U} are regular, because the only small such separation of D is (\emptyset, D) , which is in fact trivial since $|D| \geq 2$. And they are all robust; see Footnote 7. Since every $d \in D$ induces a profile $P_d := \{(A, B) \mid d \in B\}$ of U (just choose k large enough), any tree set of separations that distinguishes all the profiles in \vec{U} will contain, for every pair of distinct data $a, b \in D$, a bipartition $\{A, B\}$ of D such that $a \in A$ and $b \in B$.

Let us call the robust and regular profiles in U the *clusters* of D . Theorem 5 then becomes a consequence of Theorem 3.6, via Corollary 3.8:

Theorem 5. Every submodular data set has a canonical regular tree set of separations which efficiently distinguishes all its distinguishable clusters. \square

5.2 Edge-tangles and Gomory-Hu trees

In this section we illustrate Theorem 3.9 by applying it to some specific profiles, sometimes called ‘edge-tangles’. We reobtain a classical theorem of Gomory and Hu [16], which asserts the existence of certain trees displaying optimal cuts in a graph. Our point is not that obtaining these trees via Theorem 3.9 is simpler than their elementary existence proof, although that is not exactly short either. But seeing what Theorem 3.9 does in this special case can illustrate its power: it shows, in a familiar setting, what it can do more generally with profiles that arise in other contexts, such as [8].

Consider a finite graph $G = (V, E)$ with a positive real-valued weight function g on the edges. A tree T on V , not necessarily a subgraph of G , is a *Gomory-Hu tree* for G and g if for all distinct $u, v \in V$ there is an edge on the u - v path in T whose fundamental cut in G has minimum g -weight among all u - v cuts of G . (For details, see Frank [13, Ch. 7.2.2].)

Theorem 5.1. (Gomory & Hu, 1961)

Every finite graph with positive real edge weights has a Gomory-Hu tree.

Proof. To deduce this theorem from Theorem 3.9, consider the universe \vec{U} of all set separations of V that are bipartitions of V with the order function $|A, B| := \sum \{g(ab) \mid a \in A, b \in B, ab \in E(G)\}$, which is clearly submodular. Every vertex $v \in V$ induces a robust regular profile $P_v := \{(A, B) \mid v \in B\}$ of U , and $v \mapsto P_v$ is a bijection from V to the set \mathcal{P} of all these profiles of U .

Let $T' \subseteq U$ be the tree set which Theorem 3.9 finds for \mathcal{P} . For every $v \in V$, the profile P_v of U induces a profile t_v of $T' \subseteq U$, and $v \mapsto P_v \mapsto t_v$ is a bijection from V to the set \mathcal{O} of all the profiles of T' . (Since T' is a tree set, its profiles are just its consistent orientations.)

Every $s \in T'$ induces a bipartition of \mathcal{O} into the set $\mathcal{O}_{\vec{s}}$ of profiles that orient s as \vec{s} and the set $\mathcal{O}_{\overleftarrow{s}}$ of profiles that orient s as \overleftarrow{s} . And s is itself a bipartition of V , say $\vec{s} = (U, W)$. These bipartitions correspond via our bijection $v \mapsto t_v$:

$$U = \{v \in V \mid \overleftarrow{s} \in t_v\} = \{v \in V \mid t_v \in \mathcal{O}_{\overleftarrow{s}}\}$$

$$W = \{v \in V \mid \vec{s} \in t_v\} = \{v \in V \mid t_v \in \mathcal{O}_{\vec{s}}\}.$$

In other words, our bijection $v \mapsto t_v$ maps $\vec{s} = (U, W)$ to $(\mathcal{O}_{\overleftarrow{s}}, \mathcal{O}_{\vec{s}})$.

It is not hard to show [7] that for every $s \in T'$ there exist unique $t_u, t_v \in \mathcal{O}$ that differ only on s . Then T' becomes the edge set of a graph \mathcal{T} on \mathcal{O} if we let s join t_u to t_v , for every $s \in T'$. This graph \mathcal{T} is in fact a tree [7]. For every edge s of \mathcal{T} , the vertex sets $\mathcal{O}_{\vec{s}}, \mathcal{O}_{\overleftarrow{s}}$ induce the two components of the forest $\mathcal{T} - s$, because profiles adjacent as vertices of $\mathcal{T} - s$ agree on s .

Let T be the tree on V that is the image of \mathcal{T} under our bijection $t_v \mapsto v$ from \mathcal{O} to V . Every edge e of T then arises from an edge s of \mathcal{T} , and the bipartition $\{U, W\}$ of V defined by the components of $T - e$ corresponds via $v \mapsto t_v$ to the bipartition $\{\mathcal{O}_{\overleftarrow{s}}, \mathcal{O}_{\vec{s}}\}$ of \mathcal{O} defined by the components of $\mathcal{T} - s$. As shown earlier, this means that $s = \{U, W\}$.

Let us show that T is a Gomory-Hu tree for G and g . For all distinct $u, v \in T$ the separations in T' that distinguish t_u from t_v (equivalently: P_u from P_v) are precisely the edges on the t_u - t_v path in \mathcal{T} . One of these, s say, distinguishes P_u from P_v efficiently, by the choice of T' . By definition of T , the edge $e \in T$ corresponding to s lies on the u - v path in T . The bipartition of V defined by the two components of $T - e$ is precisely s . The fact that s distinguishes P_u from P_v efficiently means that the cut of G it defines has minimum weight among all the u - v cuts of G , as desired. \square

6 Width duality for profiles in graphs

All our theorems deal with distinguishing existing profiles in a given universe of separations, but say nothing if such profiles do not exist. For fixed k , for example, we can ask what we can say about the structure of a graph that has no k -profile. Such graphs do exist:

Example 8. Let G be a graph with a tree-decomposition of width $< k - 1$, one whose parts have order $< k$. Any k -profile P in G will orient all the separations associated with an edge of the decomposition tree T , and will thus induce a consistent orientation of $E(T)$. This orientation will point towards some node t of T .

Now consider a separation of G associated with an edge $t't$ at t , say (A, B) , with B containing the part V_t corresponding to t . Let $A' := A \cup V_t$. Then (A', B) is still in \vec{S}_k , since $|A' \cap B| = |V_t| < k$, and $(A', B) \in P$ by Proposition 4.4.

If T has another edge at t , say $t''t$, do the same to obtain $(A'', C) \in P$. Then, by (P), also $(A' \cup A'', B \cap C) \in P$; note that this separation too has order $|V_t| < k$.

Repeating this step for all the edges of T at t , we eventually arrive at $(V, V_t) \in P$. If $|V_t| \leq k - 2$, which we can easily ensure, this contradicts the consistency of P , since (V, V_t) will then be co-trivial: just add a vertex to its V_t -side to obtain a larger separation than (V_t, V) that is still small.

For tangles, Robertson and Seymour [20] address the problem of characterizing the structure of graphs without a tangle of given order: they show that a graph either has a k -tangle or a tree-decomposition¹⁶ that is clearly incompatible with the existence of a k -tangle, and which thus witnesses its non-existence. Similar duality theorems of this type were obtained in [9] for other tangle-like structures, consistent orientations of separation systems avoiding some given set \mathcal{F} of ‘forbidden subsets’. (For classical tangles, \mathcal{F} would consist of the triples $\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\}$ satisfying (T).) All these are applications of a fundamental such duality theorem for abstract separation systems [10].

The sets \mathcal{F} for which such duality theorems are proved in [10] have to consist of ‘stars’: sets of separations all pointing towards each other. For classical tangles this does not matter: one easily shows that a consistent orientation of S_k , say, avoids all the triples as in (T) as soon as it avoids the stars satisfying (T). For profiles, however, this is not the case – which is why [9] has no duality theorem for profiles.

However, it was shown in [8] that profiles can be characterized by a forbidden set \mathcal{F} of stars of separations after all, and thus have a duality theorem. Indeed, for each k there is a set $\mathcal{T}_{\mathcal{P}}(k)$ of tree-decompositions which graphs with a k -profile cannot have, and which graphs without a k -profile must have:

Theorem 6.1. [8] *For every finite graph G and every integer $k > 0$ exactly one of the following statements holds:*

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

A similar duality theorem holds more generally for regular profiles of abstract separation systems [8, Theorem 12].

7 An open problem for tangles

Unlike blocks and tangles, profiles can exist in abstract separation systems that need not even consist of separations of sets. This motivates the notion of a profile independently of their examples in graphs or matroids.

But for profiles in graphs, it is meaningful to ask just how different arbitrary profiles can be from block profiles: are they always induced by some fixed set of vertices in the same way as a block profile $P_k(b)$ is induced by b ?

This problem appears to be open even for tangles:

Problem. Given a k -tangle θ in a graph G , is there always a set X of vertices such that a separation (A, B) of order $< k$ lies in θ if and only if $|A \cap X| < |B \cap X|$?

¹⁶The ‘branch-decompositions’ used in [20] are easily translated into tree-decompositions [9].

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, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs I. Existence and algorithms. *J. Combin. Theory Ser. B*, 116:1–24, 2016.
- [3] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs II. Essential parts. *J. Combin. Theory Ser. B*, 118:268–283, 2016.
- [4] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. Connectivity and tree structure in finite graphs. *Combinatorica*, 34(1):1–35, 2014.
- [5] R. Diestel. *Graph Theory* (5th edition). Springer-Verlag, 2017. Electronic edition available at <http://diestel-graph-theory.com/>.
- [6] R. Diestel. Abstract separation systems. *Order*, 35:157–170, 2018.
- [7] R. Diestel. Tree sets. *Order*, 35:171–192, 2018.
- [8] R. Diestel, P. Eberenz, and J. Erde. Duality theorem for blocks and tangles in graphs; arXiv:1605.09139. *SIAM J. Discrete Math.*, 31(3):1514–1528, 2017.
- [9] R. Diestel and S. Oum. Tangle-tree duality in graphs, matroids and beyond. *Combinatorica*, 39:879–910, 2019. arXiv:1701.02651.
- [10] R. Diestel and S. Oum. Tangle-tree duality in abstract separation systems. *Advances in Mathematics (to appear)*, arXiv:1701.02509.
- [11] R. Diestel and G. Whittle. Tangles and the Mona Lisa. arXiv:1603.06652.
- [12] P. Eberenz. Characteristics of profiles. MSc dissertation, Hamburg 2015.
- [13] A. Frank. *Connections in Combinatorial Optimization*. Oxford University Press, 2011.
- [14] J. Geelen, B. Gerards, N. Robertson, and G. Whittle. Obstructions to branch-decomposition of matroids. *J. Combin. Theory (Series B)*, 96:560–570, 2006.
- [15] J. Geelen, B. Gerards, and G. Whittle. Tangles, tree-decompositions and grids in matroids. *J. Combin. Theory (Series B)*, 99:657–667, 2009.
- [16] R. E. Gomory and T. C. Hu. Multi-terminal network flows. *J. Soc. Ind. Appl. Math.*, 9:551–570, 1961.
- [17] M. Grohe. Quasi-4-connected components. arXiv:1602.04505, 2016.

- [18] F. Hundertmark. Profiles. An algebraic approach to combinatorial connectivity. arXiv:1110.6207, 2011.
- [19] J. Oxley. *Matroid Theory*. Oxford University Press, 1992.
- [20] N. Robertson and P. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)*, 52:153–190, 1991.
- [21] W. T. Tutte. *Graph Theory*. Addison-Wesley, 1984.