

Deciders for tangles of set separations

Reinhard Diestel Christian Elbracht Raphael W. Jacobs

Abstract

Tangles, as introduced by Robertson and Seymour, were designed as an indirect way of capturing clusters in graphs and matroids. They have since been shown to capture clusters in much broader discrete structures too. But not all tangles are induced by a cluster. We characterise those that are.

We offer two such characterisations. The first is in terms of how many small sides of a tangle's separations it takes to cover the graph or matroid. The second uses a new notion of duality for separation systems that do not necessarily come from graphs or matroids.

1 Introduction

Tangles were introduced by Robertson and Seymour as a tool in their graph minors project [15]. They provided a novel, indirect, way to capture highly cohesive substructures, or ‘clusters’, in graphs. The idea is that since clusters cannot be divided into significantly large parts by graph separations of low order, any given cluster implicitly orients every low-order separation towards its ‘big’ side, the side that contains most of the cluster. It turned out that this induced orientation of all the low-order separations collectively contains all the information needed to prove fundamental theorems about the cluster structure of a graph, which has made tangles a powerful tool in the connectivity theory of graphs.

Over the last decade, the notion of tangles has been significantly generalised to other discrete structures. These include matroids, but also bespoke structures that come with concrete clustering applications [1, 7, 8, 9, 10, 12, 13]. This has been made possible by re-casting tangle theory in terms of a purely algebraic framework of ‘abstract separation systems’, which encompass the notions of separation from all these various different contexts. Although these abstract separation systems are very general, the central tangle theorems are still valid in this framework.

In this paper we use the notation of abstract separation systems for improved readability, but in fact consider almost exclusively tangles of set separations. These include graph separations, and it may be helpful for the reader to keep the example of graph separations in mind throughout this introduction.

Suppose we are given a set S of *separations* of some set V , a set of unordered pairs $s = \{A, B\}$ of subsets of V such that $A \cup B = V$. Each of these has two *orientations*: the orientation (A, B) , which we think of as pointing towards B , and its *inverse* (B, A) . We usually denote the orientations of s by arrows on s : one of them is denoted as \vec{s} , the other as \overleftarrow{s} , but it does not matter which is which.

Now consider a subset $X \subseteq V$ that is a concrete ‘cluster’ in the sense that the separations in S cannot divide it evenly, because for every $\{A, B\} \in S$ more

than two thirds (say) of X lies in $A \setminus B$ or in $B \setminus A$. If most of X lies in $B \setminus A$, say, then X *orients* this separation *towards* B , as (A, B) .

Any orientation τ of (all the separations in) S induced by a cluster X in this way has the following property, which no longer refers to X : whenever τ orients three separations $\{A_i, B_i\} \in S$ ($i = 1, 2, 3$) towards B_i , their ‘small sides’ A_i cannot cover V , because each contains less than a third of X . This, essentially, is the definition of a *tangle* of S : any orientation of S such that no three small sides cover V . (Note that the meaning of ‘small’ here is intrinsic to τ : the side of a separation towards which τ orients it is now called ‘big’, its other side ‘small’.)

This definition of a tangle has made it possible to investigate clusters in a graph or data set without referring to them directly in the usual concrete way, as sets of vertices or data points. In particular, one can investigate the relative structure of clusters without even having found them in this concrete sense – a sense which, moreover, may be inadequate given the fuzzy nature of many real-world clusters.

However, tangles are not equivalent to clusters but weaker: while every cluster, including fuzzy clusters, gives rise to a tangle, not every tangle is induced by a cluster. Tangles that do not come from clusters can still be interesting; the *text tangles* in [4, 6] are a typical example in the context of set separations.

To be a little more formal, let us say that a set $X \subseteq V$ *decides* an orientation τ of S if, for every $(A, B) \in \tau$, there are more elements of X in B than in A . It is an open question whether every tangle of a finite graph is decided by some set of its vertices.¹ Tangles of more general set separations, however, need not have such deciders. Let us construct a simple example.²

The basic idea of our construction is that we start with S and an ‘orientation’ $\tau = \{\vec{s} \mid s \in S\}$ of S as just a collection of names, and then construct V by assigning its elements directly to the sides of the separations in τ to make τ into a tangle that has no decider set.

To implement this formally, pick an integer $\ell > 0$, and assign to each of the sets $\sigma = \{\vec{s}_1, \vec{s}_2, \vec{s}_3\}$ with $s_1, s_2, s_3 \in S$ distinct an ℓ -set V_σ so that these sets V_σ are disjoint for different σ , and let V be their union. Put the elements of V_σ in B_i for each of the three $\vec{s}_i = (A_i, B_i) \in \sigma$. Then for every $s \in S$ and $\vec{s} = (A, B)$, say, we have $B = \bigcup\{V_\sigma \mid \vec{s} \in \sigma\}$ and $A = V \setminus B$. If $|S| = m$, say, then $|V| = \binom{m}{3}\ell$, and for every $\vec{s} = (A, B) \in \vec{S}$ we have $|B| = \binom{m-1}{2}\ell$. Let us choose S so that $m \geq 6$.

By construction, the forward orientations \vec{s} of all the $s \in S$ form a tangle τ of S : the large sides of any three $\vec{s} \in \tau$ meet in $\ell > 0$ elements. Moreover,

$$\text{every } v \in V \text{ lies on the big side of exactly three elements of } \tau. \quad (*)$$

A simple double count now shows that τ has no decider set $X \subseteq V$. Indeed, suppose it does and let

$$d := \sum_{(A,B) \in \tau} |X \cap B|.$$

¹It was shown in [11] that graph tangles have *weighted* deciders, in which every vertex is counted with some weight assigned to it. These weights depend on the tangle, but not on the separation to be decided.

²The separation system S in this example will not be submodular. But it is shown in [11] that we can find a submodular order function on the set of bipartitions of V such that the separations in τ are precisely the maximal separations (in the usual partial order on the oriented separations) of a k -tangle for a suitably chosen $k \in \mathbb{N}$.

Then $|X \cap B| > |X|/2$ for every $(A, B) \in \tau$, and hence $d > m|X|/2$. On the other hand, by (*), each $x \in X$ lies in B for exactly three $(A, B) \in \tau$, so d counts it three times: $d = 3|X|$. Putting these together we obtain $3|X| > m|X|/2$. This implies $m < 6$, contrary to our assumption.

Is it possible to distil from this example some property of τ that identifies *all* the tangles without a decider set? We grappled with this question for quite a while, until we found the following solution.

Given an integer k , we say that an orientation τ of S is *k-resilient* if it takes more than k elements of τ to obtain V as the union of their small sides. Every tangle, by definition, is 3-resilient.

Our earlier example of a tangle τ without a decider set is not 4-resilient. In fact, given any four separations in τ , by (*) every $v \in V$ lies on the small side of at least one of them, so the four small sides have union V . At the other extreme, every *principal* tangle of a set of bipartitions, one consisting of all its separations (A, B) whose big side B contains some fixed element x , is *infinitely resilient* in that it is k -resilient for every $k \in \mathbb{N}$. Note that $\{x\}$ is a decider set for this tangle. In Section 3 we shall see that the unique 5-tangle of the $(n \times n)$ -grid, which is decided by its entire vertex set, is $\Omega(n^2)$ -resilient.

These examples seem to suggest that tangles of set separations that are k -resilient for large k are more likely to have decider sets. We can indeed prove such a fact, with an interesting additional twist: ‘large’ has to be measured not in terms of $|V|$ or $|S|$, but relative to the number of maximal elements of the tangle in the usual partial order of oriented separations.

This dependence on the number of maximal elements in a tangle is not even surprising: in a k -resilient tangle with at most k maximal elements, the intersection of all their big sides is non-empty, and is clearly a decider set for this tangle. If we allow weighted decider sets, or *deciders* for short, we even get a sharp bound. Roughly speaking, we can show the following (Theorem 2):

A tangle with m maximal elements has a decider if it is k -resilient for some $k > \frac{m}{2}$, which is best possible as a general bound for all tangles.

To prove this theorem we introduce the notion of *local decidability* which generalises the idea of resilience. We show in Theorem 6 that an orientation τ of a system S of set separations has a decider if and only if there exist suitable parameters k and ℓ such that τ is k -locally ℓ -decidable. In particular, there are such suitable parameters k and ℓ if an orientation τ of S is highly resilient compared with its number of maximal elements, our Theorem 2.

For our second characterisation of tangles with deciders, we exploit the recent notion of *duality* of separation systems. Dual separation systems naturally arise in applications of tangle theory. For example, consider an online shop with a set V of items on sale and a history P of purchases made last year [2, 4]. In this setting, two different separation systems occur. Every purchase in P induces a bipartition of V into the items bought versus those not bought. Equally, every item in V defines a bipartition of P into those purchases that included it versus those that did not. The tangles of these two separation systems can be shown to interact [5], and they will help us to obtain a second characterisation of tangles with deciders, Theorem 7. This will also imply Theorem 2.

Our third contribution in this paper shows the existence of decider *sets* for some tangles in the universe of set separations with the order of a separation

$\{A, B\}$ defined as $|A \cap B|$. Elbracht, Kneip, and Teegen [11] showed that the k -tangles in this universe have weighted decider sets. In Section 5 we strengthen this by showing that if a k -tangle in this universe extends to a $2k$ -tangle, it has a decider set without weights. It would be interesting to know whether similar results hold in the universe of set separations with arbitrary submodular order functions.

2 Preliminaries

This section collects together the definitions we need in this paper. While we shall work only with set separations as considered in the introduction, we use the notation and wording of the more general framework of ‘abstract separation systems’ [1]. For the convenience of those readers not familiar with this framework, we recall in this section all definitions around abstract separation systems that we shall need.

2.1 Set separations and abstract separation systems

A *separation* of a set V is a set $\{A, B\}$ consisting of two subsets A and B of V which satisfy $A \cup B = V$. The two sets A and B are the *sides* of the separation; their intersection $A \cap B$ is its *separator*. The *order* of a set separation is the size of its separator, unless specified otherwise. If the separator is empty, the separation is a *bipartition* of V . Bipartitions have order 0 with respect to the above order function for set separations, so they are usually equipped with other order functions. See [2] for various examples.

The two ordered pairs (A, B) and (B, A) are the *orientations* of $\{A, B\}$. Given such an *oriented separation* (A, B) , we call A its *small side* and B its *big side*. For two oriented separations of V we let $(A, B) \leq (C, D)$ if $A \subseteq C$ and $B \supseteq D$. Note that if (A, B) and (C, D) are bipartitions, then $A \subseteq C$ is equivalent to $B \supseteq D$. If we map every separation (A, B) to its *inverse* (B, A) , then this map forms an involution which is order-reversing in that

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

Sets of oriented set separations closed under this involution form an instance of so-called ‘abstract separation systems’, defined as follows. A *separation system* is a poset (\vec{S}, \leq) together with an order-reversing involution $*$. The elements of \vec{S} are called *oriented separations*. Given an element $\vec{s} \in \vec{S}$, its *inverse* \vec{s}^* is denoted by \overleftarrow{s} and vice-versa. As the map $*$ is order-reversing, we thus have $\vec{r} \leq \vec{s} \iff \overleftarrow{r} \geq \overleftarrow{s}$.

An *unoriented separation* s is a set of the form $\{\vec{s}, \overleftarrow{s}\}$ for some oriented separation \vec{s} . Then \vec{s} and \overleftarrow{s} are the two *orientations* of s . The set of all separations which have an orientation in a separation system \vec{S} is denoted as S .

An oriented separation \vec{s} is *trivial* in \vec{S} if there exists a $t \in S$ with $t \neq s$ such that both $\vec{s} \leq \vec{t}$ and $\overleftarrow{s} \leq \overleftarrow{t}$; its inverse \overleftarrow{s} is called *co-trivial*. We call an oriented separation \vec{s} *small* (and \overleftarrow{s} *co-small*) if $\vec{s} \leq \overleftarrow{s}$. In particular, every trivial separation is small. A separation s is *degenerate* if $\vec{s} = \overleftarrow{s}$. Finally, a *star* is a set σ of oriented, nondegenerate separations such that $\vec{r} \leq \overleftarrow{s}$ for every two distinct $\vec{r}, \vec{s} \in \sigma$.

The only small separations of a set V are of the form (A, V) , and the only degenerate one is (V, V) . Hence, (\emptyset, V) is the only small bipartition of V , and it is also trivial in every separation system S containing more than one unoriented separation.

We write any supremum of two separations \vec{r}, \vec{s} in a given separation system as $\vec{r} \vee \vec{s}$, and their infimum as $\vec{r} \wedge \vec{s}$. A separation system \vec{U} in which these always exist, i.e., which is a lattice, is a *universe* of separations and formally denoted as $(\vec{U}, \leq, *, \wedge, \vee)$. Since $*$ is order-reversing, the supremum and the infimum satisfy DeMorgan's law, in that $(\vec{r} \wedge \vec{s})^* = \vec{r} \vee \vec{s}$. In the universe of all separations of a set V , the infimum of $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$ is $(A \cap C, B \cup D)$, their supremum $(A \cup C, B \cap D)$.

Given a universe \vec{U} of separations, a map $|\cdot| : \vec{U} \rightarrow \mathbb{Z}$ is an *order function* if $|\vec{s}| = |\vec{s}^*| = |s| \geq 0$ for all $\vec{s} \in \vec{U}$. Given $k \in \mathbb{N}$, we then write

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

The order function, and the universe \vec{U} itself, are called *submodular* if

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

for all $\vec{r}, \vec{s} \in \vec{U}$. The universe of all separations of a given set, with the order of a separation $\{A, B\}$ defined as $|A \cap B|$, is a *standard universe of set separations*.

In a submodular universe of separations, each \vec{S}_k has the following property: if $\vec{r}, \vec{s} \in \vec{S}_k$, then at least one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ is also in \vec{S}_k . We say that a separation system \vec{S} inside a (not necessarily submodular) universe \vec{U} of separations is (*structurally*) *submodular* if it has this property, i.e., if for every two $\vec{r}, \vec{s} \in \vec{S}$ at least one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ also lies in \vec{S} .

Another example of a submodular universe of set separations is given by the separations of a graph $G = (V, E)$: a *graph separation* (A, B) of G is a set separation of V such that G has no edges between $A \setminus B$ and $B \setminus A$.

2.2 Orientations of separation systems

Let \vec{S} be a separation system. Assigning to every $s \in S$ either \vec{s} or \vec{s}^* is called *orienting* S (or the $s \in S$). So an *orientation* of \vec{S} or S is a set $\tau \subseteq \vec{S}$ with $|\tau \cap \{\vec{s}, \vec{s}^*\}| = 1$ for every $s \in S$. An orientation τ' of $S' \subseteq S$ *extends* to an orientation τ of S if $\tau \cap S' = \tau'$. An orientation τ of S is *consistent* if for every $\vec{r} \in \vec{S}$ such that $\vec{r} \leq \vec{s} \in \tau$ for some $s \neq r$ we also have $\vec{r} \in \tau$.

Assume now that \vec{S} lies in some universe \vec{U} of separations. An orientation τ of S is a *profile* if it is consistent and has the *profile property* in that for $\vec{s}, \vec{t} \in \tau$ we never have $(\vec{s} \vee \vec{t})^* \in \tau$. Such a profile is *regular* if it contains no co-small separation. If \vec{U} is equipped with an order function $s \mapsto |s|$, then a profile of $\vec{S}_k = \{ \vec{s} \in \vec{U} : |s| < k \}$ is a *k-profile in \vec{U}* .

Let $\mathcal{F} \subseteq 2^{\vec{U}}$. An orientation τ of S is an \mathcal{F} -*tangle* if it is consistent and *avoids* the set \mathcal{F} in that $\sigma \not\subseteq \tau$ for every $\sigma \in \mathcal{F}$. The \mathcal{T} -*tangles* of S , where

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

are also called (*abstract*) *tangles* of S ; we often just call them *tangles* of S . So the choice of $\mathcal{F} = \mathcal{T}$ is our default choice for \mathcal{F} , and we will make it clear explicitly if we consider \mathcal{F} -tangles for a different choice of \mathcal{F} .

Abstract tangles of S are examples of profiles of S [7]. In a universe \vec{U} of set bipartitions, another important class is formed by the \mathcal{F}_ℓ -tangles, where

$$\mathcal{F}_\ell = \left\{ F \subseteq \vec{U} : |F| \leq 3 \text{ and } \left| \bigcap_{(A,B) \in F} B \right| < \ell \right\} \text{ for } \ell \in \mathbb{N}.$$

2.3 Weight functions and deciders

A *weight function* on a finite set V is a map w from V to $\mathbb{R}_{\geq 0}$. For subsets $U \subseteq V$ we write $w(U) = \sum_{v \in U} w(v)$. We say that a weight function w is *non-zero* if there exists $v \in V$ with $w(v) > 0$. Note that $w(B) - w(A) = w(B \setminus A) - w(A \setminus B)$ for every set separation $\{A, B\}$ of V , a fact we shall use freely throughout. A weight function w on V with values in $\{0, 1\}$ can equivalently be formulated as an indicator function of the set $X = X_w = w^{-1}(1)$ in that $w(A) = |X \cap A|$ for every $A \subseteq V$. We shall also use this equivalence freely throughout.

Let w be a weight function on a set V . We say that w *decides* a separation $s = \{A, B\}$ of V if s has an orientation $\vec{s} = (A, B)$ such that $w(A) < w(B)$; we then also say that w *decides s as \vec{s}* . If S consists of separations of V , then w *decides S* if it decides each of its elements, thus defining an orientation τ of S .

Conversely, if τ is some set of oriented separations and S is its underlying set of unoriented separations, then we say that w *decides S like τ* , that it *witnesses τ* and all its elements, and call w a *decider for τ* . If w takes values in $\{0, 1\}$, we call $X = w^{-1}(1)$ a *decider set* for τ . If there exists a decider (set) for τ , then we say that τ *has a decider (set)*.

Let us note some basic observations about deciders. First observe that we can *scale* a weight function w on V by a positive scalar $\lambda > 0$ without changing the sign of $w(B) - w(A)$ for any set separation $\{A, B\}$ of V . In particular, if an orientation τ of a separation system S has a decider, then there exists a decider for τ which decides every separation in S like τ and at least with difference K for any given $K > 0$. This is because we can just scale a decider w for τ appropriately, i.e. by a factor $\lambda \geq K / (\min_{(A,B) \in \tau} (w(B) - w(A)))$.

This fact directly implies that, if a tangle τ has a decider, there also exists a weight function w witnessing τ which takes values in \mathbb{N} instead of \mathbb{R} . Indeed, suppose that w decides every separation of S like τ , with difference at least $\epsilon > 0$. Since \mathbb{Q} is dense in \mathbb{R} , we can replace $w(v) \in \mathbb{R}$ with a rational number $w'(v)$ such that $|w(v) - w'(v)| < \epsilon/|V|$. The resulting weight function w' clearly still witnesses τ . Now an appropriate scaling of w' yields the desired decider for τ taking values in \mathbb{N} .

A weight function witnesses an orientation of S as soon as it witnesses its maximal elements in the partial order on \vec{S} . We include a proof of this observation from [11] for the reader's convenience.

Lemma 1. *Let $w : V \rightarrow \mathbb{R}_{\geq 0}$ be a weight function on a set V . Let (A, B) and (C, D) be separations of V with $(C, D) \leq (A, B)$. If w witnesses (A, B) , it also witnesses (C, D) .*

Proof. Since $(C, D) \leq (A, B)$, we have $C \subseteq A$ and $D \supseteq B$. So $w(C) \leq w(A)$ and $w(D) \geq w(B)$, as w is a weight function. As w witnesses (A, B) , we have

$$w(C) \leq w(A) < w(B) \leq w(D),$$

so w witnesses (C, D) as well. □

3 Deciders and resilience

In this section we use the novel notion of *resilience* to prove a sufficient criterion for an orientation of some set S of set separations to have a decider. After that, we further generalise the concept of resilience towards the notion of *k-local ℓ -decidability* which allows us to give a characterisation of orientations of S with deciders. We begin this section by giving all the definitions around the concept of resilience.

Let \vec{S} be any abstract separation system in some universe \vec{U} of separations, and $k \in \mathbb{N}$. An orientation τ of S is *k-resilient* if no set of $\leq k$ elements of τ has a co-small supremum in \vec{U} . Note that a *k-resilient* orientation of S is also *k'-resilient* for every $k' < k$.

For example, if S is a set of separations of a set V , then τ is *k-resilient* if and only if for all sets $\sigma \subseteq \tau$ of at most size k , we have that $\bigcup \{A \mid (A, B) \in \sigma\} \neq V$ because a set separation is co-small if and only if it has the form (V, X) for some $X \subseteq V$. If S is even a set of bipartitions of V , then this is equivalent to $\bigcap \{B \mid (A, B) \in \sigma\} \neq \emptyset$ for all sets $\sigma \subseteq \tau$ of size at most k because (V, \emptyset) is the only co-small bipartition of V .

We can similarly define κ -resilience for an infinite cardinal κ : we call an orientation τ *κ -resilient* if no set of $\leq \kappa$ elements of τ has a co-small supremum in \vec{U} . If κ is a limit cardinal, we shall moreover say that τ is *κ^- -resilient* if τ is α -resilient for all cardinals $\alpha < \kappa$. As a special case we say that an orientation τ is *infinitely resilient* if τ is \aleph_0^- -resilient, i.e. if τ is *k-resilient* for all $k \in \mathbb{N}$.

If for some orientation τ there exists a maximal cardinal κ such that τ is κ -resilient, then we call κ the *resilience* of τ . Similarly, if there is no such maximal cardinal, but a minimal cardinal κ such that τ is not κ -resilient, then we define the *resilience* of τ to be κ^- .

In addition to the examples on resilience in the introduction, let us here illustrate the concept once more with a less extreme example. Consider the tangle τ of the set S of all graph separations of order at most 4 of the $(n \times n)$ -grid. This tangle τ has the entire vertex set of the grid as its decider set. Let us show that τ is $\Omega(n^2)$ -resilient. Notice that every element (A, B) of τ satisfies $|A| \leq 10$; indeed, most satisfy $|A| \leq 5$. Since all co-small separations of V are of the form (V, X) for some $X \subseteq V$, any set of separations in \vec{S} with a co-small supremum has at least $n^2/10$ elements.

Why can the notion of resilience help us with constructing a decider for a given orientation of set separations? Consider an orientation τ of a set S of separations of some finite ground set V . Write $M = M(\tau)$ for the set of maximal elements of τ in the partial order on \vec{S} . Let us see how high resilience of τ compared with $|M|$ might help us build a decider for τ .

Assume that τ is *k-resilient* for some integer k , and write \mathcal{M} for the set of all *k*-element subsets of $M = M(\tau)$. Then for every $M' \in \mathcal{M}$, there exists an element $v_{M'}$ of our ground set V which is strictly on the big side of all separations in M' . It seems natural to construct a decider for M (and thus for τ by Lemma 1) by combining all these local decider sets $\{v_{M'}\}$, i.e. by assigning to each $v \in V$ as its weight the number of sets $M' \in \mathcal{M}$ with $v_{M'} = v$.

It turns out that the weight function w defined in this way need not in general be a decider for M . This is because each $v_{M'}$, while adding its weight to the correct sides of the separations in M' (the sides selected by τ) can also

add weight to the wrong sides of separations in $M \setminus M'$, the sides not selected by τ . But as soon as k is big enough that each fixed separation $(A, B) \in M$ is contained in the majority of the sets in \mathcal{M} , which will happen as soon as M has more $(k - 1)$ -subsets to form a k -subset with (A, B) than it has k -subsets not including (A, B) , the orientation (A, B) of $\{A, B\}$ will be witnessed by the majority of the local decider sets $\{v_{M'}\}$ for $M' \in \mathcal{M}$. We can then deduct from this that w is a decider for M , and hence for τ .

More precisely, we have the following theorem.

Theorem 2. *Let \vec{U} be the universe of all separations of some finite set V , and let τ be an orientation of some set $S \subseteq U$ of separations. Let m be the number of maximal elements of τ in the partial order on \vec{U} . If τ is k -resilient for some $k > \frac{m}{2}$, then τ has a decider.*

We will formally obtain Theorem 2 below as a corollary of the more general Theorem 6. But before we do so, let us first show that there exists a tangle without a decider which in some sense witnesses the optimality of Theorem 2; we can find such a tangle by using a more general version of the construction from the introduction.

Proposition 3. *For all $m, k \in \mathbb{N}$ with $3 \leq k \leq \frac{m}{2}$, there exists a submodular universe \vec{U} of set bipartitions and an m -tangle $\tau_{m,k}$ in \vec{U} which has m maximal elements in \vec{U} and is k -resilient, but which does not have a decider.*

An example of such tangles is given by a certain type of hypergraph edge tangles introduced in [11]. Let us describe the construction of these tangles $\tau_{m,k}$ given natural numbers $m \geq k \geq 3$, and then we show that the tangles $\tau_{m,k}$ do indeed have all the desired properties.

Proof of Proposition 3. We consider the universe \vec{U} of all bipartitions of the set V which consists of all k -element subsets of $[m] = \{1, \dots, m\}$. We set $V_i = \{X \in V : i \in X\}$ for every $i \in [m]$, and equip U with the order function $|\cdot|$ which assigns to $\{A, B\} \in U$ the order

$$|\{A, B\}| = |\bigcup A \cap \bigcup B| = m - |\{i \in [m] : V_i \subseteq A \text{ or } V_i \subseteq B\}|.$$

This order function is easily seen to be submodular (see [11] for a formal proof).

For each $i \in [m]$, let $\vec{s}_i = (V \setminus V_i, V_i) \in \vec{U}$. Clearly, $|s_i| = m - 1$. We claim that $M = \{\vec{s}_1, \dots, \vec{s}_m\}$ is the set of maximal elements of an m -tangle in \vec{U} which we shall denote by $\tau_{m,k}$. To see this, let us first show that every separation $r = \{A, B\} \in U$ of order $< m$ has an orientation \vec{r} with $\vec{r} \leq \vec{s}_i$ for some $i \in [m]$. By the definition of the \vec{s}_i , it is enough to show that A or B must include V_i for some $i \in [m]$. For if not, we have $V_i \cap A \neq \emptyset$ and $V_i \cap B \neq \emptyset$ for all $i \in [m]$. In particular, we have $i \in \bigcup A \cap \bigcup B$ for every $i \in [m]$, and thus $|\{A, B\}| = m$ which is a contradiction. Thus, M is the set of maximal elements of some consistent orientation $\tau_{m,k}$ of $S_m^{\vec{}}$. To see that $\tau_{m,k}$ is indeed an m -tangle, it is enough to observe that $\vec{s}_i \vee \vec{s}_j \vee \vec{s}_l$ is not co-small for any $1 \leq i, j, l \leq m$. This is immediate since there exists an element of V contained in the proper right sides of all three of $\vec{s}_i, \vec{s}_j, \vec{s}_l$ as $k \geq 3$.

As shown in [11], it is immediate from double counting (as in the example from the introduction) that the m -tangle $\tau_{m,k}$ does not have a decider for $m \geq 6$

if $k \leq \frac{m}{2}$. So it remains to show that $\tau_{m,k}$ is indeed k -resilient. But this is immediate from the construction: by Lemma 1, it is enough to consider an arbitrary collection $s_{i_1}^{\rightarrow}, \dots, s_{i_k}^{\rightarrow}$ of k separations in the set M of maximal separations of $\tau_{m,k}$. Then the set $\{i_1, \dots, i_k\} \in V$ is on the big side V_{i_j} of $s_{i_j}^{\rightarrow}$ for all $j \in [k]$ by construction, and thus the union of the small sides of the bipartitions $s_{i_j}^{\rightarrow}$ does not contain $\{i_1, \dots, i_k\}$. \square

The $\tau_{m,k}$ constructed in Proposition 3 are abstract tangles, but the construction can easily be modified to find \mathcal{F}_ℓ -tangles for arbitrary $\ell \geq 2$ with the same properties: instead of taking the k -element subsets of $[m]$ as our ground set V , we can take as V the disjoint union of ℓ -element sets V_σ , one for each k -element subset σ of $[m]$. Then the example from the introduction for a tangle without a decider set reappears here as the \mathcal{F}_ℓ -tangle form of $\tau_{6,3}$.

Before we proceed towards a proof of Theorem 2, let us briefly investigate our examples of \mathcal{F}_ℓ -tangles without deciders in some more detail. Note that the construction of \mathcal{F}_ℓ -tangles as above works for constant ℓ only: it does not necessarily work when the value of ℓ is not constant, but large in terms of $|V|$, e.g. of size at least $\epsilon|V|$ for some constant $\epsilon > 0$.

The following proposition shows that there exists a sharp lower bound for those $\epsilon > 0$ for which $\ell \geq \epsilon|V|$ guarantees the existence of decider. Let $\tau_{m,k}$ be the tangle constructed in the proof of Proposition 3.

Proposition 4. *Let V be a set of size n , and let $0 < \epsilon < 1$. If $\epsilon \geq 1/8$, then every $\mathcal{F}_{\epsilon n}$ -tangle of a set S of bipartitions of V has a decider; if $\epsilon > 1/8$, it even has a decider set.*

Conversely, for every $\epsilon < 1/8$ there exist integers m, n and k such that $\tau_{m,k}$ is an $\mathcal{F}_{\epsilon n}$ -tangle of bipartitions of an n -set that has no decider.

Proof. Let τ be an \mathcal{F}_ℓ -tangle on a set S of bipartitions of some set V with $\ell \geq |V|/8$. If all of V is a decider set for τ , then we are done; so suppose not. Then there exists a separation $(A_1, B_1) \in \tau$ with $|B_1| \leq |V|/2$. Again we are done if B_1 is a decider set for τ . If this is not the case, then there exists a separation $(A_2, B_2) \in \tau$ such that $|B_1 \cap A_2| \geq |B_1 \cap B_2|$; in particular, we have $|B_1 \cap B_2| \leq |V|/4$.

It turns out that if $\ell > |V|/8$, then $B_1 \cap B_2$ needs to be the desired decider set since otherwise, there exists another separation $(A_3, B_3) \in \tau$ such that

$$|(B_1 \cap B_2) \cap B_3| \leq |(B_1 \cap B_2) \cap A_3|.$$

This implies $|(B_1 \cap B_2) \cap B_3| \leq |V|/8$ which contradicts the fact that τ is an \mathcal{F}_ℓ -tangle.

In the case of $\ell = |V|/8$, the same arguments as above result in a decider set if at least one of the occurring inequalities is strict. So suppose that all the inequalities are satisfied with equality. In particular, every separation (A, B) needs to satisfy $|A| \leq |B|$. With a similar reasoning as above, we can at least obtain a decider: note first that it is enough to find a weight function w on V which witnesses $\tau' \subseteq \tau$, where τ' consists of all separations $(A, B) \in \tau$ with $|A| = |B|$. Given a decider w for τ' , we obtain a decider for τ by adding large enough constant weight to all vertices in V .

Suppose there are two separations $(A, B), (C, D) \in \tau'$ with $|B \cap C| > |B \cap D|$. Then this yields $|B \cap D| < |V|/4$ which in turn implies the existence of a de-

cider set with the same arguments as above. Consequently, for every two separations $(A, B), (C, D) \in \tau'$, we have $|B \cap C| \leq |B \cap D|$.

Hence, the weight function w defined by counting for every $v \in V$ the number of those separations $(A, B) \in \tau'$ with $v \in B$ is a decider for τ' : given some separation $(C, D) \in \tau'$ we have that $w(C) = \sum_{(A, B) \in \tau'} |B \cap C|$ and $w(D) = \sum_{(A, B) \in \tau'} |B \cap D|$. As above, we have $|B \cap C| \leq |B \cap D|$ for every separation $(A, B) \in \tau'$; as we clearly have $|D \cap C| < |D \cap D|$, this implies $w(C) < w(D)$. Thus, w is a decider for τ .

For the second part of the proposition, let us consider the tangle $\tau_{m,k}$ as constructed in Proposition 3 for some $m \geq 2k \geq 6$. Then for any three maximal elements of $\tau_{m,k}$, their intersection contains exactly $\binom{m-3}{k-3}$ elements of the constructed ground set V . In particular, $\tau_{m,k}$ is an \mathcal{F}_ℓ -tangle for all $\ell < \binom{m-3}{k-3}$. Now recall that the size of the ground set V is $|V| = \binom{m}{k}$. Thus, if we set $k = m/2$, then $\lim_{m \rightarrow \infty} \binom{m-3}{k-3} / \binom{m}{k} = 1/8$. So by Proposition 3, we find for any $\epsilon < 1/8$ some $n \in \mathbb{N}$ and integers m, k such that the tangle $\tau_{m,k}$ witnesses that there exists an $\mathcal{F}_{\epsilon n}$ -tangle on a ground set of n elements without a decider. \square

Back to Theorem 2, recall that this theorem is sharp in terms of the parameter k in k -resilience as shown in Proposition 3. But the converse of Theorem 2 fails, i.e. not even every tangle with a decider set has high resilience.

Example 5. Let \vec{S} be the separation system consisting of all bipartitions of the set $V = [n]$ that have a side of size $< n/3$. Let τ be the orientation of \vec{S} which orients every separation s towards the side which contains more elements. In particular, V is a decider set for τ .

This orientation τ of S is a tangle, since no three big sides of separations in τ have empty intersection. However, four big sides can, so the supremum of four separations in τ can be co-small. Thus, τ has resilience 3.

Now τ has $m = \binom{n}{\lceil n/3 \rceil - 1}$ maximal elements, namely those separations whose small side has maximum size. In particular, the resilience of τ is low compared with m , although τ has a decider and even a decider set.

It turns out that we can generalise the notion of resilience in a way which includes the tangle from the previous Example 5 without invalidating Theorem 2. In fact, our more general notion leads to a more general result, Theorem 6, which actually characterises the orientations with deciders.

Our more general notion of resilience is based on our earlier observation that k -resilience provides a one-element decider set $\{v_{M'}\}$ for every k -set $M' \subseteq M$ where M is the set of maximal separations of a k -resilient orientation τ of set separations. In the following definition we ask, instead of k -resilience, that there exist a ‘local’ decider $w_{M'}$ for every k -set $M' \subseteq M$ which decides M' correctly and simultaneously is not too badly wrong on the separations in $M \setminus M'$.

More precisely, we call an orientation τ of separations of some finite set V *k -locally ℓ -decidable* for given $k \in \mathbb{N}$ and $\ell \geq 0$ if for every set $M' \subseteq \tau$ of size $|M'| \leq k$, there is a weight function $w_{M'}$ on V such that

- (i) $\forall (A, B) \in M' : w_{M'}(B) - w_{M'}(A) \geq 1$;
- (ii) $\forall (A, B) \in \tau : w_{M'}(A) - w_{M'}(B) \leq \ell$.

Observe that in the above definition one need only consider sets M' of separations that are maximal in τ by Lemma 1. In addition, we can equivalently strengthen the condition of $|M'| \leq k$ to $|M'| = k$, as long as τ has at least k maximal elements.

Note that the above definition is indeed a generalisation of our earlier notion of resilience since, if τ is k -resilient, then it is k -locally 1-decidable with $w_{M'}$ assigning 1 to a single element in $V \setminus \bigcup_{(A,B) \in M'} A$ and 0 to all the other elements in V .

Let us show that, if τ has a decider w , then τ is k -locally ℓ -decidable for all $k \in \mathbb{N}$ and $\ell \geq 0$. Any decider for τ clearly satisfies (ii). In (i), a decider w for τ only guarantees > 0 rather than ≥ 1 as required. But since τ is finite, we can obtain the latter by scaling w suitably. In particular, the tangle τ in Example 5 is k -locally ℓ -decidable for every $k \in \mathbb{N}$ and $\ell \geq 0$.

Here, then, is our generalisation of Theorem 2:

Theorem 6. *Let \vec{U} be the universe of all separations of some finite set V , and let τ be an orientation of some set $S \subseteq U$. Let M be the set of maximal elements of τ in the partial order on \vec{U} , and $m = |M|$. Then τ has a decider if and only if it is k -locally ℓ -decidable for some $k \in \mathbb{N}$ and $\ell > 0$ with $k > \frac{m}{1+1/\ell}$.*

Since, as noted earlier, every k -resilient orientation τ is k -locally 1-decidable, Theorem 2 is a direct corollary of Theorem 6.

Proof of Theorem 6. If τ has a decider w , then it is by definition k -locally ℓ -decidable for every $k \in \mathbb{N}$ and $\ell \geq 0$. In particular, it is m -locally ℓ -decidable for every $\ell > 0$, and we have $m > \frac{m}{1+1/\ell}$ in this case.

For the converse, it is enough to show by Lemma 1 that M has a decider. We note that if $k \geq m$ the statement is true immediately, so suppose for the following that $k < m$. We construct a decider w for M as follows: write \mathcal{M} for the set of all k -element subsets of $M = M(\tau)$. For every $M' \in \mathcal{M}$, we have a weight function $w_{M'}$ as in the definition of k -local ℓ -decidability. Then we combine all these weight functions to define

$$w : V \rightarrow \mathbb{R}_{\geq 0}, w(v) = \sum_{M' \in \mathcal{M}} w_{M'}(v).$$

We show that w is the desired decider for M .

Consider $(A, B) \in M$. Then (A, B) is contained in $\binom{m-1}{k-1}$ sets $M' \in \mathcal{M}$, and similarly (A, B) is not contained in $\binom{m-1}{k}$ such $M' \in \mathcal{M}$. So by the definition of k -locally ℓ -decidable, we get by (i) that

$$\sum_{(A,B) \in M'} w_{M'}(A \setminus B) \leq \sum_{(A,B) \notin M'} w_{M'}(B \setminus A) + \ell \cdot \binom{m-1}{k},$$

writing $(A, B) \in M'$ as a shortcut to mean that we sum over all those sets $M' \in \mathcal{M}$ containing (A, B) . Similarly, we obtain by (ii) that

$$\sum_{(A,B) \notin M'} w_{M'}(A \setminus B) \leq \sum_{(A,B) \in M'} w_{M'}(B \setminus A) - \binom{m-1}{k-1}.$$

These inequalities combine to

$$\begin{aligned}
w(A \setminus B) &= \sum_{M' \in \mathcal{M}} w_{M'}(A \setminus B) \\
&= \sum_{(A,B) \notin M'} w_{M'}(A \setminus B) + \sum_{(A,B) \in M'} w_{M'}(A \setminus B) \\
&\leq \sum_{(A,B) \notin M'} w_{M'}(B \setminus A) + \ell \cdot \binom{m-1}{k} + \sum_{(A,B) \in M'} w_{M'}(B \setminus A) - \binom{m-1}{k-1}.
\end{aligned}$$

Now since $k > \frac{m}{1+1/\ell}$ and $k < m$, we have that $\ell < \frac{k}{m-k}$ and thus

$$\ell \cdot \binom{m-1}{k} < \binom{m-1}{k-1},$$

as $\binom{m-1}{k} = \frac{(m-1)!}{k!(m-1-k)!}$ and $\binom{m-1}{k-1} = \frac{(m-1)!}{(k-1)!(m-k)!}$ differ precisely in the factor $\frac{k}{m-k}$. This implies

$$\begin{aligned}
w(A \setminus B) &< \sum_{(A,B) \notin M'} w_{M'}(B \setminus A) + \sum_{(A,B) \in M'} w_{M'}(B \setminus A) \\
&= \sum_{M' \in \mathcal{M}} w_{M'}(B \setminus A) = w(B \setminus A).
\end{aligned}$$

So w witnesses M and hence τ . □

4 Deciders and duality

In this section we present a second characterisation of the orientations of set separations that have a decider: one in terms of a duality between systems of set separations. As an unexpected corollary, we obtain an independent second proof of Theorem 2.

The duality of set separations, which was introduced in [2] and first studied in [3, 5], is defined as follows. Given a system \vec{S} of separations of a set V , we start by picking for every $s \in S$ a default orientation, which we denote as \vec{s} (rather than \bar{s}). If we think of $\vec{s} = (A, B)$ as the side $B \subseteq V$ to which it points (so that \bar{s} is equated with A by the same token), then for every $v \in V$ the sets

$$\vec{v} = \{s \in S \mid v \in \vec{s}\} \quad \text{and} \quad \bar{v} = \{s \in S \mid v \in \bar{s}\}$$

form a separation $\{\vec{v}, \bar{v}\}$ of the set S . Let us assume that the sets $\{\vec{v}, \bar{v}\}$ differ for distinct $v \in V$, just as the sets $\{\vec{s}, \bar{s}\}$ differ for distinct s by definition of s .³ Then they determine their v uniquely, and we may think of each v as shorthand for $\{\vec{v}, \bar{v}\}$. This makes V into a set of separations $v = \{\vec{v}, \bar{v}\}$ of S and

$$\vec{V} := \{\vec{v} \mid v \in V\} \cup \{\bar{v} \mid v \in V\}$$

³Recall that separation systems are formally defined in such a way that their elements \vec{s} are given first, and s is then formally defined as $\{\vec{s}, \bar{s}\}$. Hence, if the s are distinct, as they are here by assumption, then this means that these 2-sets are distinct.

into the set of all orientations of elements of V , and we have

$$v \in \vec{s} \Leftrightarrow s \in \vec{v} \quad \text{as well as} \quad v \in \overleftarrow{s} \Leftrightarrow s \in \overleftarrow{v} \quad (**)$$

for all the elements $\vec{s}, \overleftarrow{s}$ of \vec{S} and $\vec{v}, \overleftarrow{v}$ of \vec{V} .

There is a natural default orientation for \vec{V} in that we orient each $v \in V$ as $(\overleftarrow{v}, \vec{v})$, i.e. towards those $s \in S$ with $v \in \vec{s}$. The dual of \vec{V} with respect to this default orientation is then again \vec{S} . In simple terms, ‘dualising the dual yields the primal’ [5].

To emphasise the role of the default orientation of \vec{S} for constructing a dual separation system, we shall use the following notation in this section: given an orientation σ of \vec{S} , we denote the dual separation system by $\vec{V} = \vec{V}(\sigma)$. The natural default orientation of \vec{V} described above is then denoted by $\tau = \tau(\sigma)$.

In the context of deciders, we can now ask whether the existence of a decider for an orientation σ of S relates to any property of the natural default orientation $\tau = \tau(\sigma)$ of the dual $V(\sigma)$ of S . It turns out that it does, and it does so in an intriguing way: σ has a decider if and only if every non-zero weight function w' on the ground set S of $V(\sigma)$ witnesses the orientation of some separation in τ .

Theorem 7. *Let S be a set of separations of some finite set V , and let σ be any orientation of S . Write $V(\sigma)$ for the dual separation system of S with respect to σ , and let $\tau = \tau(\sigma)$ be the default orientation of $V(\sigma)$. Then the following two assertions are equivalent:*

- (i) *There exists a decider for σ .*
- (ii) *For every non-zero weight function w' on S , there exists $\vec{v} = (A', B') \in \tau$ with*

$$w'(A') = \sum_{s \notin \vec{v}} w'(s) < \sum_{s \in \vec{v}} w'(s) = w'(B').$$

As $s \mapsto \vec{s}$ is a bijection between S and σ , we could equivalently define the weight function w' in Theorem 7 (ii) on the orientation σ of S ; for notational simplicity we will freely switch between these two definitions.

The proof of Theorem 7 will be done in terms of pure linear algebra and can formally be followed without any further knowledge about the duality of set separations. However, it may be fruitful for the reader to think of the objects used in this proof from a homological point of view on duality as introduced in [3]. Let us briefly sketch the idea behind this homological perspective.

We consider the free abelian groups C_0 and C_1 with bases V and S , respectively. The elements of C_0 are called the *0-chains* of our separation system S , and the elements of C_1 are its *1-chains*. Then the *boundary operator* $\partial : C_1 \rightarrow C_0$ is defined by sending the 1-chain $s = (A, B) \in S$ to the 0-chain $\sum_{v \in B \setminus A} v - \sum_{v \in A \setminus B} v$. The matrix representation Q of this boundary homomorphism is precisely the matrix that we are going to consider in the subsequent proof of Theorem 7.

Now the transpose Q^T of this matrix is the matrix representation of the ‘co-boundary operator’: for $n \in \{0, 1\}$, let $C^n = \text{Hom}(C_n, \mathbb{Z})$ be the set of n -cochains of our separation system S . The *co-boundary operator* $\delta : C^0 \rightarrow C^1$ then sends a 0-cochain $\phi : C_0 \rightarrow \mathbb{Z}$ to the 1-cochain $(\phi \circ \partial) : C_1 \rightarrow \mathbb{Z}$. So

the duality of set separations as described in the beginning of this section is precisely reflected in the duality of the boundary operator ∂ and the co-boundary operator δ . For a detailed discussion about this homological perspective on duality, we refer the reader to [5].

Now an orientation σ of S forms a basis of C_1 and similarly, its natural dual orientation τ of V forms a basis of C_0 . So given a weight function w on τ (or equivalently on V), we can identify it with a vector $x = w(\tau) \in \mathbb{R}_{\geq 0}^n$ where $n = |V|$. The co-boundary operator Q^T maps x onto a homomorphism from C_1 to $\mathbb{R}_{\geq 0}$ (we here implicitly extend Q^T to real coefficients). By definition, this homomorphism $Q^T x$ takes a positive value for all $\vec{s} \in \sigma$ if and only if w is a decider for σ (see Theorem 7 (i)). Analogously, we can derive a similar statement (see Theorem 7 (ii)) for the boundary operator Q applied to $y = w'(\sigma) \in \mathbb{R}_{\geq 0}^l$ for a weight function w' on σ (or equivalently on S). These properties of the (co-)boundary operators are precisely the ones that we are making use of in the proof of Theorem 7 when considering the matrix Q .

Let us now prove Theorem 7. The key tool in this proof will be the following variant of Farkas' Lemma (see e.g. [14, 6. Theorem]⁴).

Lemma 8 (Farkas' Lemma). *Let $Q \in \mathbb{R}^{n \times l}$ and $b \in \mathbb{R}^l$. Then exactly one of the following two assertions holds:*

- (i) *There exists $x \in \mathbb{R}_{\geq 0}^n$ with $Q^T x \geq b$.*
- (ii) *There exists $y \in \mathbb{R}_{\geq 0}^l$ with $Qy \leq 0$ and $b^T y > 0$.*

Proof of Theorem 7. Fix enumerations $V = \{v_1, \dots, v_n\}$ and $\sigma = \{\vec{s}_1, \dots, \vec{s}_l\}$ writing $\vec{s}_i = (A_i, B_i)$ for $i \in [l]$. Using these enumerations, we shall, for the course of this proof, identify a weight function w on V with a vector $x = x(w) \in \mathbb{R}_{\geq 0}^n$ and a weight function w' on S with a vector $y = y(w') \in \mathbb{R}_{\geq 0}^l$.

Let us define a matrix $Q = Q(\sigma) \in \mathbb{R}^{n \times l}$ via

$$Q_{ij} = \begin{cases} 1, & v_i \in B_j \setminus A_j; \\ 0, & v_i \in A_j \cap B_j; \\ -1, & v_i \in A_j \setminus B_j. \end{cases}$$

For each $s \in S$, let \vec{s} denote the orientation of s in σ . Analogously, let \vec{v} denote the orientation of v in τ .

Now given a weight function w' on S , we obtain for $y = y(w') \in \mathbb{R}_{\geq 0}^l$ that

$$(Qy)_i = \sum_{v_i \in B_j} y_j - \sum_{v_i \in A_j} y_j = \sum_{s \in \vec{v}_i} w'(s) - \sum_{s \notin \vec{v}_i} w'(s) \quad \forall i \in [n].$$

So a non-zero weight function w' on S is *not* as in Theorem 7 (ii) if and only if $Qy \leq 0$ for $y = y(w') \in \mathbb{R}_{\geq 0}^l$ where \leq is meant coordinate-wise.

Similarly, let w be a weight function on V and $x = x(w) \in \mathbb{R}_{\geq 0}^n$. Then we compute

$$(Q^T x)_j = \sum_{v_i \in B_j} x_i - \sum_{v_i \in A_j} x_i = \sum_{v \in \vec{s}_j} w(v) - \sum_{v \notin \vec{s}_j} w(v) \quad \forall j \in [l].$$

⁴Our version of Farkas' Lemma follows from [14, 6. Theorem] by applying their theorem to $A = Q^T$ and $-b$ instead of b .

So a weight function w on V is a decider for σ if and only if $Q^T x > 0$ for $x = x(w) \in \mathbb{R}_{\geq 0}^n$. Recall from Section 2 that a decider for σ can be scaled arbitrarily by a positive scalar and still witnesses σ . Hence, there exists a decider for σ (as in Theorem 7 (i)) if and only if $Q^T x \geq \mathbf{1}$ for $x = x(w) \in \mathbb{R}_{\geq 0}^n$ and some non-zero weight function w on V , where $\mathbf{1}$ denotes the constant 1 vector.

The result then follows by applying Lemma 8 to Q and $b = \mathbf{1} \in \mathbb{R}^l$, and denoting $x = x(w)$ in Lemma 8 (i) and $y = y(w')$ in Lemma 8 (ii). \square

By Lemma 1, an orientation σ has a decider if and only if its set of maximal elements has a decider. Therefore, we have the following corollary of Theorem 7 with almost the same proof.

Corollary 9. *Let S be a system of separations of some finite set V . Let σ be an orientation of S , and let $M = M(\sigma)$ be the set of maximal elements of σ . Then the following two assertions are equivalent:*

- (i) *There exists a decider for σ .*
- (ii) *For every non-zero weight function w' on M , there exists $v \in V$ with*

$$\sum_{\vec{s} \in \{\vec{r} \in M \mid v \notin \vec{r}\}} w'(\vec{s}) < \sum_{\vec{s} \in \{\vec{r} \in M \mid v \in \vec{r}\}} w'(\vec{s}). \quad \square$$

As an illustration of the power of Theorem 7, let us re-prove Theorem 2 about the existence of deciders for highly resilient orientations.

Proposition 10. *Let \vec{U} be the universe of all separations of some finite set V , and let σ be an orientation of a set $S \subseteq U$ of separations. Let $M = M(\sigma)$ be the set of maximal elements of σ , and write $m = |M|$. If σ is k -resilient for some $k > \frac{m}{2}$, then σ has a decider.*

Proof. We apply Corollary 9 in that we consider an arbitrary non-zero weight function w' on M and show that Case (ii) in Corollary 9 holds. Let $M' \subseteq M$ consist of those k separations in M which have the highest weight with respect to w' . Since σ is k -resilient, there exists some $v \in V$ which lies strictly on the big side of $\bigvee M'$. Now this v is as desired, since

$$\begin{aligned} \sum_{\vec{s} \in \{\vec{r} \in M \mid v \in \vec{r}\}} w'(\vec{s}) &\geq \sum_{\vec{s} \in M'} w'(\vec{s}) > \frac{1}{2} \sum_{\vec{s} \in M} w'(\vec{s}) \\ &\geq \sum_{\vec{s} \notin M'} w'(\vec{s}) \geq \sum_{\vec{s} \in \{\vec{r} \in M \mid v \notin \vec{r}\}} w'(\vec{s}). \end{aligned}$$

\square

5 Deciders for extendable tangles

Let \vec{U} be a submodular universe of separations of a ground set V , and let τ be an orientation of S_k for some $k \in \mathbb{N}$. In Section 3 we analysed different properties of τ which ensure the existence of a decider for τ . All these properties required

us to consider large subsets of τ instead of the usual triples which are required for the definition of a tangle. In particular, all the notions considered above may be viewed as a strengthening of the triple condition in the definition of a tangle, i.e. we give a stronger condition that an orientation needs to satisfy in order to be a tangle with a decider.

But how can we guarantee the existence of a decider for a tangle τ if we do not want to strengthen the definition of a tangle in the above sense? We know that there exists tangles without deciders (see e.g. Proposition 3). So instead of looking for a decider for τ itself, we may try to find a decider for some subset τ' of the separations of τ . Ideally, we can do so in such a way that this decider for the subset is still, in some sense, related to the original tangle τ .

In the presence of an order function, one natural such subset of a k -tangle τ , say, consists of all separations of order less than some $k' < k$. In other words, we would like to obtain, given a k -tangle τ , a decider for the k' -tangle $\tau' \subseteq \tau$. One way in which we could try to achieve this consists in proving the following: if a tangle τ' extends to some tangle τ of higher order in \vec{U} , then τ' has a decider. In this case, we may view the decider w for τ' as an approximation of a decider for its extension τ – although w will in general not decide all the separations in τ . The m -tangles $\tau_{m,k}$ constructed in Proposition 3, for example, do not have a decider; but if we consider only those separations of order at most $\frac{m}{2}$ in this example, then they even have a decider set: the whole ground set V decides all separations of order at most $\frac{m}{2}$ like $\tau_{m,k}$.

This leads us to the question of whether tangles which extend to tangles of twice their order always have deciders, or even decider sets. We now show that they do: k -profiles which extend to regular $2k$ -profiles have decider sets, as long as we work in the standard universe \vec{U} of set separations. Recall that all k -profiles in such universes have deciders [11], but those as above even have decider sets:

Theorem 11. *Let \vec{U} be the standard universe of separations of a finite set V . If τ' is a k -profile in \vec{U} that extends to a regular $2k$ -profile τ in \vec{U} , then τ' has a decider set $X \subseteq V$ of size $|X| \geq 2k$.*

The proof of Theorem 11 will find a star $\sigma \subseteq \tau$ whose interior $\bigcap_{(A,B) \in \sigma} B$ is the desired decider set for τ' . Let us show first that the interior of any star in τ has size at least $2k$.

Lemma 12. *The interior of any star in a regular $2k$ -profile in \vec{U} has at least $2k$ elements.*

Proof. Suppose not, let τ be a regular $2k$ -profile in \vec{U} , and let $\sigma \subseteq \tau$ be a star whose interior $X = \bigcap_{(A,B) \in \sigma} B$ has size $|X| < 2k$.

Let us write $\sigma = \{(A_1, B_1), \dots, (A_l, B_l)\}$. We claim that for any $i \leq l$ we have $|(A_1, B_1) \vee \dots \vee (A_i, B_i)| < 2k$. By definition, we have

$$|(A_1, B_1) \vee \dots \vee (A_i, B_i)| = |(A_1 \cup \dots \cup A_i) \cap (B_1 \cap \dots \cap B_i)|.$$

Since σ is a star, we have $(A_1 \cup \dots \cup A_i) \subseteq B_j$ for every $j > i$. So in particular, we have

$$(A_1 \cup \dots \cup A_i) \cap (B_1 \cap \dots \cap B_i) \subseteq (B_{i+1} \cap \dots \cap B_l) \cap (B_1 \cap \dots \cap B_i) = X.$$

Therefore, $|(A_1, B_1) \vee \dots \vee (A_l, B_l)| \leq |X| < 2k$. By the profile property of τ , it follows inductively that $((A_1, B_1) \vee \dots \vee (A_l, B_l)) \in \tau$ for every $l \leq l$. Then the separation $(A_1, B_1) \vee \dots \vee (A_l, B_l) = (Y, X)$ is in τ where $Y = \bigcup_{(A,B) \in \sigma} A$.

Since $|X| < 2k$, the separation $\{X, V\}$ has order $< 2k$ and hence an orientation in τ . By the regularity of τ , this orientation must be (X, V) because (V, X) is co-small. But this leads to a contradiction since this would imply $(Y, X) \vee (X, V) = (V, X) \in \tau$ by the profile property of τ . \square

Proof of Theorem 11. Let σ be a star in τ with an interior $X = \bigcap_{(A,B) \in \sigma} B$ of smallest possible size. By Lemma 12 we have $|X| \geq 2k$. We claim that X is the desired decider set for τ' .

For this suppose that X does not witness $(A, B) \in \tau'$. Since $|X| \geq 2k$, we then especially have $|X \cap A| \geq k$ which we are going to lead to a contradiction.

So let $(A, B) \in \tau'$ be of minimal order among all separations with $|X \cap A| \geq k$. Note that this separation (A, B) may be witnessed by X . For every $(C, D) \in \sigma$, the corner separation $(A \cap D, B \cup C)$ has at least the order of (A, B) as otherwise $(A \cap D, B \cup C)$ would contradict the choice of (A, B) : indeed, by construction, we have $X \subseteq D$, and therefore $|(A \cap D) \cap X| = |A \cap X| \geq k$. Thus, by the minimality of $|(A, B)|$, the corner $(A \cap D, B \cup C)$ must have order at least $|(A, B)|$.

By submodularity, the opposite corner $(B \cap C, A \cup D)$ has order at most $|(C, D)|$ and thus we have $(B \cap C, A \cup D) \in \tau$ by consistency. Now consider the star $\hat{\sigma} \subseteq \tau$ consisting of (A, B) together with, for every $(C, D) \in \sigma$, the separation $(B \cap C, A \cup D)$.

We claim that the interior \hat{X} of $\hat{\sigma}$ is smaller than X contradicting the choice of σ . Indeed, by definition, we have

$$\hat{X} = B \cap \bigcap_{(C,D) \in \sigma} (A \cup D) = (A \cap B) \cup (B \cap X) = ((A \cap B) \setminus X) \cup (B \cap X).$$

Since X is the disjoint union of $B \cap X$ and $(A \cap X) \setminus B$, we are done if

$$|(A \cap B) \setminus X| < |(A \cap X) \setminus B|.$$

Let $h = |A \cap B \cap X|$. Since $|A \cap X| \geq k$, we have $|(A \cap X) \setminus B| \geq k - h$. However, we have $(A, B) \in \tau$, so $|A \cap B| < k$ and hence

$$|(A \cap B) \setminus X| = |A \cap B| - |A \cap B \cap X| < k - h$$

completing the proof. \square

Our proof of Theorem 11 heavily relies on the assumption that the order function on \vec{U} is given by $|(A, B)| = |A \cap B|$. We do not know whether a similar result holds for other or even all submodular order functions on such \vec{U} .

Problem 13. *Let \vec{U} be the universe of all separations of a finite set V , equipped with any submodular order function. Is it true that if τ' is a k -profile in \vec{U} which extends to a regular $2k$ -profile in \vec{U} then τ has a decider set X ? What happens in other universes of set separations, such as the universe of bipartitions of V ?*

References

- [1] R. Diestel. Abstract separation systems. *Order*, 35:157–170, 2018.
- [2] R. Diestel. Tangles in the social sciences: a new mathematical model to identify types and predict behaviour. arXiv:1907.07341, 2019.
- [3] R. Diestel. Hypergraph homology. arXiv:2007.09125, 2020.
- [4] R. Diestel. *Tangles: indirect clustering in the empirical sciences*. In preparation.
- [5] R. Diestel, J. Erde, C. Elbracht, and M. Teegen. Duality and tangles of set partitions. In preparation.
- [6] R. Diestel, J. Erde, and K. Krishnareddy. Tangles: a novel way to classify texts. In preparation.
- [7] R. Diestel, J. Erde, and D. Weißauer. Structural submodularity and tangles in abstract separation systems. *J. Combin. Theory (Series A)*, 167C:155–180, 2019.
- [8] R. Diestel, F. Hundertmark, and S. Lemanczyk. Profiles of separations: in graphs, matroids, and beyond. *Combinatorica*, 39(1):37–75, 2019.
- [9] R. Diestel and S. Oum. Tangle-tree duality in abstract separation systems. *Advances in Mathematics*, 377:107470, 2021.
- [10] C. Elbracht and J. Kneip. A canonical tree-of-tangles theorem for structurally submodular separation systems. arXiv:2011.09758, to appear in *Combinatorial Theory*, 2020.
- [11] C. Elbracht, J. Kneip, and M. Teegen. Tangles are decided by weighted vertex sets. *Advances in Comb.*, 2020:9, 2020.
- [12] C. Elbracht, J. Kneip, and M. Teegen. Trees of tangles in abstract separation systems. *J. Combin. Theory (Series A)*, 180:105425, 2021.
- [13] J. Geelen, B. Gerards, N. Robertson, and G. Whittle. Obstructions to branch-decomposition of matroids. *J. Combin. Theory (Series B)*, 96:560–570, 2006.
- [14] C. Perng. On a class of theorems equivalent to Farkas’s lemma. *Applied Mathematical Sciences*, 11:2175–2184, 2017.
- [15] N. Robertson and P. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)*, 52:153–190, 1991.