Tangles, Trees of Tangles, and Submodularity

DISSERTATION

zur Erlangung der Würde des
Doktors der Naturwissenschaften

der Fakultät für Mathematik, Informatik
und Naturwissenschaften,

Fachbereich Mathematik, der

UNIVERSITÄT HAMBURG

vorgelegt von

Maximilian Teegen

aus Hamburg.

Hamburg, 2021

Vorsitzender der Prüfungskommission:
Prof. Armin Iske (Universität Hamburg)

Erstgutachter und Betreuer:
Prof. Reinhard Diestel (Universität Hamburg)

Zweitgutachter:
Joshua Erde, PhD (Technische Universität Graz)

Datum der Disputation:
2. September 2021

Contents

1. Introduction

This thesis concerns itself with abstract separation systems and their tangles, both in general and in some concrete instances.

Background

Tangles were originate from graph minor theory, they were invented as a tool in the monumental graph minors project of Robertson and Seymour [64]. In the original sense, a tangle describes a well-connected or 'dense' region in a graph indirectly: not by specifying a subset of vertices or edges, but by declaring, for each way of cutting the graph into two parts along few vertices, which side the dense region lies on. Formally, a *separation* of a graph G is an ordered pair (A, B) of subsets of $V(G)$ where $A \cup B = V$ and no edge of G runs between $A \setminus B$ and $B \setminus A$. The size of the *separator* $A \cap B$ is *order* of that separation. A *k-tangle* τ in G is a set that contains for every separation (A, B) of G of order less than k exactly one of (A, B) and (B, A) and fulfils the following consistency condition:

$$\forall (A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau \colon G[A_1] \cup G[A_2] \cup G[A_3] \neq G. \tag{T}$$

We think of a separation (A, B) in a tangle as pointing 'away from A, towards B'. This consistency condition (T) ensures that the directions that a tangle points in are not contradictory to each other.

Tangles were used to study how tree-like the structure of a graph is in terms of *tree-decompositions* and *tree-width*. Informally, the former is a way of building up our graph from small subgraphs, glued to each other in the shape of a tree; the latter measures with how small subgraphs we can achieve such a decomposition: the lower the tree-width, the smaller the parts, and the more tree-like is our graph. Tree-width is within a constant factor of *branch-width*, a measure based on another kind of decomposition along a tree, and low branch-width is dual to the existence of high order tangles: a graph has branch-width $< k$ if and only if it has no tangle of order k. That statement is the *tangle–tree-duality theorem* of Robertson and Seymour [64].

Accompanying the duality theorem as the second foundation of tangle-theory is what we call the *tree-of-tangles theorem* – also by Robertson and Seymour [64] – which (roughly speaking) says that multiple different tangles in a graph – even across different orders k – are arranged in relation to each other in a tree: there is a tree-decomposition, such that each tangle points towards a different part.

Since the inception of tangles in [64], the general idea has been transferred to various other concepts of clusters, width, and tree-like decompositions, for graphs as well as other discrete structures – primarily matroids – see e.g. [22, 25–27, 29, 33, 44, 46, 48, 50, 61].

As part of this expansion of the tangle concept, and in an effort to simplify and axiomatize the theory, generalizing it to cover many of these varations, one particular branch of tangle theory evolved: the theory of *abstract separation systems*. These separation systems are a generalization established in the work of Diestel and Oum [28] and of Diestel, Hundertmark, and Lemanczyk [26], and they proved to be a very general and useful concept which generalizes the notion of a set of separations. The vertex separations of a graph introduced above and separations of matroids both are instances of this abstract framework.

Formally, a *separation system* $(\vec{S}, \leqslant, ^*)$ is a partially ordered set whose elements are called *separations*, together with an order-reversing involution *. The partial order $\leqslant$ represents the specificity of a separation, for separations (A, B) and (C, D) of a graph this means that $(A, B) \leqslant (C, D)$ whenever $A \subseteq C$ and $B \supseteq D$. The involution * 'inverts' a separation, for graph separations it is the function $(A, B) \mapsto (B, A)$.

Tangles will then be *orientations* of such a separation system – chosing for each separation it or its inverse – which satisfies some consistency condition. The various notions of 'tangle' are expressed each by a set $\mathcal{F}$ of subsets of a separation system $\vec{S}$ that are to be avoided – like the triples of separations in property (T). Orienting all separations in $\vec{S}$ in a way that includes none of these forbidden subsets then defines an *$\mathcal{F}$-tangle of $\vec{S}$*. We will make these definitions more precise in Chapter 2.

To allow an abstract representations of 'k-tangles' across multiple values of k, separation systems usually come with some more structure in the form of *universes of separations*, separation system whose partial order forms a lattice, so any pair of separations has a supremum and an infimum. The generalized expression of the order of a separation is then a *submodular order functions* defined on a universe of separations. The set of all separations of a graph form such a universe of separations – one then considers tangles on the subsets of separations up to some order k.

For $\mathcal{F}$-tangles, with reasonable assumption of $\mathcal{F}$, Diestel and Oum proved the abstract tangle–tree duality theorem [28]. Effectively, it still asserts that the absence of a tangle is equivalent to the existence of a tree-like decomposition – but, in the abstract setting, the concept of 'tree-like decomposition' hinges on which definitions of 'separation' and 'tangle' are used.

For the tree-of-tangles theorem, a variety of different extensions and abstractions of this theorem exist [10, 11, 13, 14, 25, 26]. Multiple variations of the tree-of-tangles theorem exist for abstract separation systems. In fact, the term is used to describe a collection of theorems, all of which state some variation of the following:

Let $\mathcal{T}$ be a (sufficiently nice) set of distinguishable tangles, then there exists a nested set of separations $N(\mathcal{T})$ that contains, for every pair of tangles in $\mathcal{T}$, a separation which distinguishes that pair of tangles.

The pairwise nestedness of separations means (up to technicalities) that they are arranged in relation to each other like the edges of a tree. The tangles that they distinguish then each point to a node of the eponymous tree. Depending on the conditions on the given set of tangles $\mathcal{T}$, the theorems guarantee additional properties of the set $N(\mathcal{T})$: canonicity, meaning invariance under isomorphisms, and/or the 'efficiency' of the distinguishing separations. We will go into detail on those later; in fact, Part II of this thesis will be dedicated entirely to tree-of-tangles theorems, discussing their requirements and how to prove them.

If one now wants to have such theorems in a concrete setting, then one only has to perform the translation between abstract separation systems and concrete separations of the setting and check that the conditions of the theorems hold. Examples of such applications can be found in [26–28].

Overview of this thesis

After a formal introduction to the terminology of abstract separation systems and tangles in Chapter 2, we present our own results. The thesis is split into three parts: Part I is about the structure that a single tangle can induce – in two concrete instances, Part II is devoted to trees of tangles and tree-of-tangles theorems, and Part III is concerned with a technical aspect of the separation systems we define our tangles on: submodularity properties and how they affect the structure of abstract separation systems.

Part I: Representing a single tangle indirectly

Part I is about the structure that a single tangle can induce. The most well-known instance of such a phenomenon is the grid theorem of Robertson and Seymour [64, (7.5)] which says that, for any integer k, every tangle τ of sufficiently high order $N(k)$ in a graph is witnessed by a $(k \times k)$-grid minor. This grid minor, in turn, induces a tangle τ' where, of every separation in the tangle, the side that the tangle points to contains the majority of the minor's branch sets. This tangle τ' is, in fact, the restriction of τ to separations of order less than k. Picking one vertex from every branch set thus gives a *decider set* for τ': a set X of vertices so that for every separation $(A, B) \in \tau'$ more vertices of X are in B than in A.

However, the large gap between $N(k)$ and k means that we cannot obtain a decider set for every tangle in this way. In [26] Diestel thus posed the following question: does every tangle in a graph have a decider set? — In Chapter 3, we show that a weighted version of this is true, i.e., that there is a non-negative weight function defined on the vertices such that, for every separation (A, B) in the tangle, the sum of the weights in B is larger than the sum of the weights in A.

⇝ p. 20 **Theorem 1.** *Let $G = (V, E)$ be a finite graph and τ a k-tangle in G. Then there exists a function $w\colon V \to \mathbb{N}$ such that a separation (A, B) of G of order $< k$ lies in τ if and only if $w(A) < w(B)$, where $w(U) := \sum_{u \in U} w(u)$ for $U \subseteq V$.*

We prove the existence of such a function even in a more general case of profiles of set separations with order function $(A, B) \mapsto |A \cap B|$, and we prove a variant of this theorem also for edge tangles, a kind of tangles defined in terms of graph cuts rather than the usual (vertex) separations.

Chapter 4 covers dual tangles in the setting of bipartite graphs. Given a bipartite graph with partition classes X and Y, we introduce a measure of how connected the sides of a separation of X are in terms of the neighbourhoods of vertices in Y. Symmetrically, we can apply the same measure of connectivity to the separations of Y. This measure plays nicely with a natural way of separations of X to induce separations of Y and vice versa: the idea is that a separation (A, B) of X induces a separation of Y by asking for every vertex in Y whether it has the majority of its neighbours in A or in B and separating the vertices in Y accordingly. Reasonable definitions of tangles on X and on Y then allow us to 'shift' tangles back-and-forth between X and Y.

Concretely, we can then show that a tangle of order $4k$ on X shifts to a tangle of order k on Y and that this operation stabilizes in the sense that the double-shift of a tangle is a subset of the original.

Part II: Trees of tangles

Chapter 5 covers the 'splinter lemmas', a framework which allows us to prove in a unified way the most relevant existing tree-of-tangles theorems. There are different variations of splinter lemmas covering non-canonical and canonical, finite and infinite settings. The underlying principle however is always the same: they all state in some way that if we can, given any two non-nested separations each distinguishing its own pair of tangles, resolve this non-nestedness locally by replacing one of the two separations with one which serves the same role and is better nested. Then applying these resolutions in the correct way globally to all pairs leads to a nested set of separations. The base version of the splinter lemma reads as follows:

⇝ p. 64 **Lemma 13** (Splinter Lemma)**.** *Let U be a universe of separations and $\mathfrak{A} = (\mathcal{A}_i)_{i \leqslant n}$ a family of subsets of U. If $\mathfrak{A}$ splinters, then we can pick an element a_i from each $\mathcal{A}_i$ so that $\{ a_1, \dots, a_n \}$ is nested.*

We use our splinter lemmas to re-prove many of the existing tree-of-tangles theorems – with and without efficiency, canonical and non-canonical – but they also allow us to introduce some new tree-of-tangles theorems of our own; for example, for profiles in sequences of separation systems:

Theorem 14. *If $\mathcal{S} = (S_1, \dots, S_n)$ is a compatible sequence of structurally submodular separation systems inside a universe U, and $\mathcal{P}$ is a robust set of profiles in $\mathcal{S}$, then there is a nested set N of separations in U which efficiently distinguishes all the distinguishable profiles in $\mathcal{P}$.* $\rightsquigarrow$ *p. 72*

Lemma 13 can be phrased in such a general way, that it can be applied in contexts which do not fit the – already very general – framework of abstract separation systems: as we will see in Section 5.6, Lemma 13 can be phrased in terms of a relation. With a bit more effort than for separation systems, it can then be applied in new settings and we present an instance of that in Section 5.7 where we consider directed separations and tangles in directed graphs.

We extend the approach of the splinter lemmas to an infinite setting in Section 5.8 with the 'thin splinter lemma', Lemma 22. This lemma allows us in Section 5.9 to give simplified proofs of two tree-of-tangles theorems for infinite graphs. The first is by Carmesin [10]:

Theorem 5.20 ([10, Theorem 5.12])**.** *For any graph G, there is a nested set N of separations that distinguishes efficiently any two robust principal profiles (that are not restrictions of one another).* $\rightsquigarrow$ *p. 97*

The other theorem is by Carmesin, Hamann, and Miraftab [14] and a deal more technical, introducing 'trees of tree-decompositions' to handle the problems that come up when constructing tree-decompositions for infinite graphs. We use Lemma 22 to establish a theorem which sits half-way between those two theorems, and from which either can be deduced more easily:

Theorem 24. *Given a set of distinguishable robust regular profiles $\mathcal{P}$ of a graph G there exists a canonical nested set of separators efficiently distinguishing any pair of profiles in $\mathcal{P}$.* $\rightsquigarrow$ *p. 109*

We present another application of Lemma 22 in Section 5.10. This is again an application to infinite graphs, but this time we are not distinguishing tangles or profiles but *edge-blocks*:

Theorem 25. *Every connected graph G has a nested set of bonds that efficiently distinguishes all the edge-blocks of G.* $\rightsquigarrow$ *p. 118*

Here, an edge-block in G is a subset of the vertices which, for some integer k, is $\subseteq$-maximal with the property that it cannot be separated in G by less than k edges. In order to prove Theorem 25 we apply Lemma 22 not to the traditional (vertex) separations of G, but rather to the separation system of all the cuts of G.

Chapter 6 also is about proving tree-of-tangles theorems. The theorems themselves are unspectacular, especially in light of Chapter 5. The extraordinary aspect is the method by which we prove those theorems: by applying the tangle–tree duality theorem. The tangle–tree duality theorem and the tree-of-tangles theorem are the two fundamental pillars on which tangle theory stands. The tangle–tree duality theorem says that if there is no tangle, then an S-tree, a tree with separations for edges, witnessing this exists. So it is rather surprising that we can utilise this in a context where many tangles do exist, the tree-of-tangles theorem. This is achieved by applying the tangle–tree duality theorem to a very carefully chosen notion of 'tangle'. In this way we push the limits of the tangle–tree duality theorem to make it produce a tree of tangles.

Finally, Chapter 7 introduces a heuristic approach to computing tangles and trees of tangles. If one wants to identify dense substructures in some large real-world dataset, say, then one might try to apply the concept of tangles. While algorithms exist for a large class of instances of tangles to compute tangles and a tree of tangles [51], in practice the size of the dataset might make that computation infeasible. This is essentially due to the fact that there are many ways to separate such a large dataset, so a way around this is to consider only a small subset of separations which are deemed 'reasonably good' according to some sampling heuristic. In Chapter 7 we explain how, with such imperfect knowledge of a separation system, one can compute tangles of the subsystem – which serve as an approximation of the tangles of the original system – and how one can combine this approach with the ideas from Chapter 5, the splinter lemmas, to build trees of tangles.

Part III: Submodularity

Submodularity is a key technical concept in tangle theory that features throughout virtually all proofs involving tangles. Originally, the theory of tangles in abstract separation was concerned only with universes $\vec{U}$ of separations which come with a *submodular order function* $|\cdot| : \vec{U} \to \mathbb{N}$. In this setting one considers the subsystems $\vec{S} = \vec{S}_k := \{ \vec{s} \in \vec{U} : |\vec{s}| < k \} \subseteq \vec{U}$ that are induced by the order function. [26, 28]

This changed when it was discovered in [25] that, for most arguments, only a structural consequence of the presence of such a submodular order function was used: that of any two separations in the system $\vec{S}$ either the supremum or the infimum (taken in the surrounding universe $\vec{U}$) lies again in $\vec{S}$. This structural condition on $\vec{S}$, called *submodularity in* $\vec{U}$, is central to much of the work in Part II.

We dedicate Part III to studying this condition of submodularity in $\vec{U}$, the stronger condition of *order-induced submodularity in* $\vec{U}$ from submodular order function, and also to a weakened, structural, condition that any two elements of $\vec{S}$ just have a supremum or an infimum in the poset $\vec{S}$ without any reference to a surrounding universe of separations.

Chapter 8 explores these three different submodularity conditions (order-induced submodularity in $\vec{U}$, submodularity in $\vec{U}$, and submodularity) in relationship to each other. We will go into further detail on these conditions and our results at the start of Chapter 8 and give just a broad overview here.

On the one hand, we will establish a link between the weakened condition, mere *submodularity* of $\vec{S}$, and the standard condition of submodularity in some universe of separations $\vec{U}$. We show that submodularity of a separation system $\vec{S}$ on its own implies submodularity in $\vec{U} \supseteq \vec{S}$ for one specific universe $\vec{U}$ that we can construct around $\vec{S}$.

On the other hand we show that order-induced submodularity in $\vec{U}$ is a proper strengthening of submodularity in $\vec{U}$, giving examples of separation systems which are submodular in some $\vec{U}$ but not order-induced submodular.

Further, we develop two decomposition theorems for submodular separation systems and, along the way, prove a handful of other interesting small facts about the different notions of submodularity.

Chapter 9 introduces one particular, fundamental question about submodular separation systems: Given a submodular separation system, is there always a separation that we can delete, so that the remainder is still submodular? A sequence of separations that we can delete one after the other while maintaining submodularity until we reach the empty set is an *unravelling*. We show the existence of unravellings for the stronger and for the weaker concept of submodularity in Section 9.2 and Section 9.4, respectively.

For the third concept, submodularity inside a fixed universe of separations, we give a counterexample in Section 9.3. In this counterexample the universe of separations is non-distributive, the distributive case remains and boils down to an intriguing-but-hard combinatorial problem which we term 'the unravelling problem':

Problem 9.1 (Unravelling problem)**.** A finite set $\mathcal{X}$ of finite sets is *woven* if, for all $X, Y \in \mathcal{X}$, at least one of $X \cup Y$ and $X \cap Y$ is in $\mathcal{X}$. Let $\mathcal{X}$ be a non-empty woven set. Does there exist an $X \in \mathcal{X}$ for which $\mathcal{X} - X$ is again woven? ⇝ *p. 189*

2. Basic Concepts

The graph-theoretic notation and terminology of this thesis follows the textbook of Diestel [19] and we assume familiarity with the general concepts and fundamental theorems of finite and infinite graph theory from there. Our terminology of tangle-theory follows the principles of abstract separation systems that evolved over the years [20, 21, 25–28]; the definitions below capture the current state of that language as used throughout this thesis, including some slightly more specific terminology of our own for subsystems and submodularity to make it precise enough to meet the needs of Chapters 8 and 9 in particular.

The definitions below are also valid for infinite posets/separation systems/graphs etc., however beyond this chapter – unless explicitly stated otherwise, as in Sections 5.8 to 5.10 – we shall assume our graphs/separation systems etc. to be finite.

2.1. Separation systems

A *separation system* $(\vec{S}, \leqslant, {}^*)$ consists of a partially ordered set $(\vec{S}, \leqslant)$ together with an involution * on $\vec{S}$ which is order-reversing.[1] That * is an *involution* on $\vec{S}$ means that it is a self-inverse function $^*\colon \vec{S} \to \vec{S}$, and that it is *order-reversing* means that if $\vec{s} \leqslant \vec{t}$ for $\vec{s}, \vec{t} \in \vec{S}$, then $\vec{s}^* \leqslant \vec{t}^*$. The elements of $\vec{S}$ are called *(oriented) separations* and the *inverse* $\vec{s}^*$ of a separation $\vec{s} \in \vec{S}$ is denoted by $\overleftarrow{s}$.

Intuitively, we think of a separation as pointing towards one of two halves of some structure, e.g., of a graph. We then understand the inverse of a separation to point towards the other half. A separation being greater than another then means, that it is more specific – pointing to a more restricted region than the other.

Given a separation $\vec{s}$, the set $s := \{\, \vec{s}, \overleftarrow{s} \,\}$ of $\vec{s}$ together with its inverse is called the corresponding *unoriented separation*. For a separation system $(\vec{S}, \leqslant, {}^*)$ we denote the set of corresponding unoriented separations as S. Conversely, given a set S of unoriented separations, we write $\vec{S}$ for the set of their orientations.

When there is no risk of confusion we use notions defined for unoriented and oriented separations interchangeably.

A *subsystem* $\vec{S}'$ of a separation system $\vec{S}$ is a subset $\vec{S}' \subseteq \vec{S}$ which is closed under involution, it inherits its partial order from $\vec{S}$.

[1]In the context of order theory separation systems are known under several different names, most notably as *involution posets* [3] or *i-posets* for short. However, so-far there has been little overlap between the aspects of separation systems that order theory is interested in and those that graph theory is interested in.

A map $\varphi\colon \vec{S} \to \vec{S}'$ is a *homomorphism of separation systems* between separation systems $\vec{S}$ and $\vec{S}'$ if it commutes with taking inverses and is order-preserving, i.e., whenever $\vec{r} \leqslant \vec{s}$ for separations $\vec{r}, \vec{s} \in \vec{S}$, then $\varphi(\vec{r}) \leqslant \varphi(\vec{s})$. An *isomorphism of separation systems*, then, is a bijection which is a homomorphism of separation systems and whose inverse is also a homomorphism of separation systems. An *embedding* is an injective homomorphism which is an isomorphism when restricting the range to the subsystem that is its image.

If for a separation $\vec{s}$ we have $\vec{s} \leqslant \overleftarrow{s}$, then we call $\vec{s}$ *small* and $\overleftarrow{s}$ *co-small*. A set of separations which does not contain any co-small separations is called *regular*.

A separation $\vec{s}$ in a separation system $\vec{S}$ is called *trivial in* $\vec{S}$ if $\vec{s} < \vec{r}$ and $\overleftarrow{s} < \vec{r}$ for some $\vec{r} \in \vec{S}$, i.e., $\vec{s} \leqslant \vec{r}$ and $\overleftarrow{s} \leqslant \vec{r}$ for some $\vec{r} \in \vec{S} \setminus s$. In this case we call $\vec{r}$ a *witness* of the triviality of $\vec{s}$ and call $\overleftarrow{s}$ *co-trivial*. Note that a trivial separation is always small, but that the converse is not generally true, since for triviality the relation to the other separations in the system is relevant. In particular, a separation which is trivial in $\vec{S}$ is not trivial in every subsystem of $\vec{S}$.

A separation $\vec{s}$ where $\overleftarrow{s} = \vec{s}$ is called *degenerate*.

2.2. Lattices and universes of separations

A *lattice* is a partially ordered set $(L, \leqslant)$ where every pair of elements $a, b \in L$ has a supremum $a \vee b$, called *join*, and an infimum $a \wedge b$, called *meet*. We usually understand a lattice as an algebraic structure $(L, \leqslant, \vee, \wedge)$, but note that the binary operators $\vee$ and $\wedge$ are completely determined by the partial order and, vice versa, the partial order is completely determined by just one of these operators, for we have

$$a \leqslant b \iff a \vee b = b \iff a \wedge b = a.$$

A *homomorphism* between lattices is a function which commutes with $\vee$ and $\wedge$ (and thus is also order-preserving). Consequently, two lattices are *isomorphic* if and only if they are isomorphic as posets. A *sublattice* of L is a subset $L' \subseteq L$ which is closed under joins and meets. It inherits the partial order of L, and thus the join and meet operations are simply the restrictions of $\vee$ and $\wedge$ from L to L'.

If a lattice has a greatest element it is called the *top* element, and we denote it as $\top$; a least element is called *bottom*, denoted $\bot$. Note that every finite lattice has a top and a bottom element. A lattice L is called *distributive* if $\vee$ and $\wedge$ are mutually distributive:

$$a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c) \;\text{ and }\; a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$$

for all $a, b, c \in L$.

The prototypical example of a distributive lattice is the *subset lattice* of some set V, which consists of the power set 2^V ordered by inclusion $\subseteq$. The join of any two subsets of V is their union, and their meet is their intersection.

A *universe of separations*[2] $(\vec{U}, \leqslant, {}^*, \vee, \wedge)$ is a separation system which is also a lattice. In a universe of separations DeMorgan's laws

$$(\vec{r} \vee \vec{s})^* = \overleftarrow{r} \wedge \overleftarrow{s} \quad \text{and} \quad (\vec{r} \wedge \vec{s})^* = \overleftarrow{r} \vee \overleftarrow{s}$$

hold. A function between universes of separations is a *homomorphism of universes of separations* if it is a homomorphism of lattices as well as a homomorphism of sepration systems.

For a pair of separations $\vec{r}$ and $\vec{s}$ in $\vec{U}$ the separations $\vec{r} \vee \vec{s}$, $\overleftarrow{r} \vee \vec{s}$, $\vec{r} \vee \overleftarrow{s}$ and $\overleftarrow{r} \vee \overleftarrow{s}$ are the *(outwards) corners of r and s in* $\vec{U}$.

2.3. Submodularity

A function $f\colon L \to \mathbb{R}$ on a lattice L is called *submodular* if it satisfies

$$f(a \vee b) + f(a \wedge b) \leqslant f(a) + f(b) \qquad \text{for all } a, b \in L.$$

If the above holds with equality always, then f is *modular*.

A function which is defined on a separation system is called *symmetric* if it is invariant under taking inverses. A *submodular order function* on a universe $\vec{U}$ then is a symmetric, non-negative, submodular function on $\vec{U}$. A universe together with a submodular order function is called a *submodular universe* and the order function is usually denoted $|\cdot|$. The symmetry allows writing $|s|$ for $|\vec{s}|$ where it is more convenient.

Given a submodular universe $(\vec{U}, \leqslant, {}^*, \vee, \wedge)$, and some number k we usually consider the subsystem

$$\vec{S}_{k,\,|\cdot|} := \{\, \vec{s} \in \vec{U} : |s| < k \,\}.$$

If the order function and universe are clear from context this is shortened to $\vec{S}_k$. For any submodular function $f\colon L \to \mathbb{R}$ on a lattice L and elements $a, b \in L$ we observe that if $f(a \vee b) > \mathsf{max}(f(a), f(b))$, then $f(a \wedge b) < \mathsf{min}(f(a), f(b))$; vice versa, if we have $f(a \wedge b) > \mathsf{max}(f(a), f(b))$, then $f(a \vee b) < \mathsf{min}(f(a), f(b))$. Hence, in the context of a submodular universe, if we have separations $\vec{r}$ and $\vec{s}$ in some $\vec{S}_k$, then at least one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ is also contained in that $\vec{S}_k$.

If, more generally, a subsystem $\vec{S}$ of a universe of separations $\vec{U}$ contains for all $\vec{r}, \vec{s} \in \vec{S}$ at least one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$, then $\vec{S}$ is called *submodular in* $\vec{U}$.[3] If, moreover, there exists a submodular order function f on $\vec{U}$ and some k such that $\vec{S} = \vec{S}_{k,\,f}$, then we say that *the submodularity of* $\vec{S}$ *in* $\vec{U}$ *is order-induced* or that $\vec{S}$

[2]In this thesis, we often write *universe* for brevity, without any intention of laying universal claim to the term.

[3]This property was originally called *structural submodularity* in [25], and the term is still common in the literature, but the independence on $\vec{U}$ in that term becomes carelessly imprecise in context of our Theorem 29.

is *order-induced submodular in* $\vec{U}$. In this case we also say that f and k *induce the submodularity* of $\vec{S}$ in $\vec{U}$.

A separation system $\vec{S}$ on its own can also be submodular: we call $\vec{S}$ *submodular*, if any two separations $\vec{s}, \vec{t} \in \vec{S}$ have either a supremum in $\vec{S}$ or an infimum in $\vec{S}$. Note that these suprema and infima ar taken in the partially ordered set $\vec{S}$ itself, not in some universe – in general, such suprema and infima need not exist, and submodularity asserts that one of them does. Every separation system that is submodular in some universe is submodular in this sense. Note that this kind of submodularity will only feature

2.4. Profiles and tangles

A set of oriented separations is called *antisymmetric* if it does not contain distinct orientations of the same separation.[4] An *orientation* of a separation system S is an antisymmetric subset of $\vec{S}$ which contains an orientation of every separation in S. In this context, an antisymmetric subset of $\vec{S}$ is also called a *partial orientation* of $\vec{S}$. A (partial) orientation O is *consistent* if there are no $\vec{r}, \vec{s} \in O$ with $\overleftarrow{r} \leqslant \vec{s}$, unless $r = s$. In other words, O is consistent if no distinct two separations in O point away from each other.

Given a separation system $\vec{S}$ and a set $\mathcal{F} \subseteq 2^{\vec{S}}$ of *forbidden subsets*, an $\mathcal{F}$*-tangle of* $\vec{S}$ is a consistent orientation of $\vec{S}$ which includes no element of $\mathcal{F}$ as a subset. When given a submodular universe and a set $\mathcal{F} \subseteq 2^{\vec{U}}$, a $\mathcal{F}$-tangle τ *in* $\vec{U}$ is a $\mathcal{F}$-tangle of $\vec{S}_k$ for some k. We then say, that τ *is of order* k. For $k < \ell$, every $\mathcal{F}$-tangle τ of $\vec{S}_\ell$ is a superset of a $\mathcal{F}$-tangle $\tau \cap \vec{S}_k$ of $\vec{S}_k$, which is called the *truncation* of τ to order k. Those $\mathcal{F}$-tangles in $\vec{U}$ which are not truncations of any other $\mathcal{F}$-tangle in $\vec{U}$ are the *maximal* $\mathcal{F}$*-tangles in* $\vec{U}$.

A consistent orientation P of a separation system $\vec{S}$ in a universe $\vec{U}$ is called a *profile of* $\vec{S}$ if it satisfies the *profile property*:

$$\forall\, \vec{r}, \vec{s} \in P\colon\ (\overleftarrow{r} \wedge \overleftarrow{s}) \notin P. \tag{P}$$

If $\vec{U}$ is a submodular universe, then a k*-profile in* $\vec{U}$ is a profile of $\vec{S}_k$ and a *profile in* $\vec{U}$ is a k-profile in $\vec{U}$ for some k. In this case P is *of order* k. Note that the profiles of a separation system $\vec{S}$ in a universe $\vec{U}$ are precisely its $\mathcal{P}$-tangles where

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

One often considers *regular profiles* only, i.e., profiles which do not contain any co-small separation. In the context of a submodular universe $\vec{U}$, a profile P in $\vec{U}$ is *robust* if the following holds:

$$\forall\, \vec{s} \in P,\ \vec{r} \in \vec{U}\colon\ (|\overleftarrow{r} \wedge \vec{s}| < |r| \text{ and } |\overleftarrow{r} \wedge \overleftarrow{s}| < |r|) \implies (\vec{r} \vee \overleftarrow{s} \in P \text{ or } \vec{r} \vee \vec{s} \in P). \tag{R}$$

[4] That is, degenerate separations are allowed to be contained in an antisymmetric set, but for any non-degenerate separation $\vec{s}$ at most one of $\vec{s}$ and $\overleftarrow{s}$ is contained.

2.5. Tree sets and S-trees

Two separations in a separation system are called *nested* if they have orientations which are comparable, that is, separations $\vec{s}$ and $\vec{t}$ are nested if $\vec{s} \leqslant \vec{t}$, $\overleftarrow{t} \leqslant \vec{s}$, $\overleftarrow{s} \leqslant \vec{t}$, or $\vec{t} \leqslant \overleftarrow{s}$. Separations which are not nested are said to *cross*. A *nested set* consists of separations which are pairwise nested.

The following elementary observation about corner separations and nestedness, the *fish lemma*, is used throughout many, if not most, proofs about tangles in abstract separation systems:

Lemma 2.1 ('Fish lemma', [20, Lemma 3.2]). *Let U be a universe of separations and $r, s \in U$ be two crossing separations. Every separation $t \in U$ that is nested with both r and s is also nested with all four corner separations of r and s.*

If the separations in a nested set form a separation system which has no trivial separations, then that separation system is called a *tree set*; and if it contains not even a small/co-small separation, then it is a regular tree-set.

Given a tree T every orientation $\vec{e} = \overrightarrow{xy} \in \vec{E}(T)$ of an edge of T defines a natural (A, B) bipartition of $V(T)$ into the component A of $T - e$ which contains the tail x and the component B of $T - e$ which contains the head y. The oriented edges of T thus naturally define a separation system $\vec{E}(T) \subseteq \vec{\mathcal{B}}(V(T))$ where we identify each edge with the corresponding bipartition of $V(T)$. This separation system is called the *edge tree set* of T.

Given a separation system $\vec{S}$, an *S-tree* (T, α) is a tree T together with a function $\alpha\colon \vec{E}(T) \to \vec{S}$ which commutes with involution. The S-tree (T, α) is *order-respecting* if α is order-preserving, i.e., $\alpha(\vec{e}) \leqslant \alpha(\vec{f})$ whenever $\vec{e} \leqslant \vec{f}$; this makes f a homomorphism of separation systems.

Every edge tree set forms a tree set and every finite regular tree set is isomorphic to the edge tree-set of a suitable tree (cf. [21] for a detailed exposition of finite tree sets). For infinite tree sets Gollin and Kneip showed the following equivalence:

Theorem 2.2 ([49, Theorem 3.9(1)]). *Let G be any connected graph, and let $\vec{S}$ be any regular tree set. Then the following assertions are equivalent:*

- *there exists an S-tree (T, α) such that $\alpha\colon \vec{E}(T) \to \vec{S}$ is an isomorphism between separation systems;*
- *$\vec{S}$ contains no chain of order-type $\omega + 1$.*

Here, the term *chain of order-type* $\omega + 1$ is meant in the usual order-theoretic sense: a set $\{\, \vec{s}_i : i \in \mathbb{N} \cup \{\omega\} \,\} \subseteq \vec{S}$ such that $\vec{s}_i < \vec{s}_j$ for all $i, j \in \mathbb{N} \cup \{\omega\}$ where $i < j$.

2.6. Important instances of separation systems and tangles

2.6.1. Set separations and vertex separations

Many important instances of separation systems consist of *set separations* of some set V: ordered pairs (A, B) of subsets of V, such that $A \cup B = V$. We denote by $\vec{S}(V)$ the universe of all set separations of V, where the involution maps every set separation (A, B) to (B, A), and where the partial order is defined by letting $(A, B) \leqslant (C, D)$ whenever $A \subseteq C$ and $B \supseteq D$. The join and meet of two separations (A, B), (C, D) in such a universe are

$$(A, B) \vee (C, D) = (A \cup C, B \cap D) \quad \text{and} \quad (A, B) \wedge (C, D) = (A \cap C, B \cup D).$$

The small separations are those of the form (X, V) for some $X \subseteq V$, and the only degenerate separation is (V, V).

For the sake of conciseness – and for historical reasons – the unoriented separation corresponding to a set separation (A, B) is usually represented by $\{A, B\}$ rather than $\{\,(A, B),\ (B, A)\,\}$.

Given a graph G, a *vertex separation* of G is a set separation (A, B) of $V(G)$ where every edge of G runs within either A or B, i.e., no edge of G runs between $A \setminus B$ and $B \setminus A$. The set $A \cap B$ is called the *separator* of such a separation, and every connected component of $G - (A \cap B)$ is contained fully in either $A \setminus B$ or $B \setminus A$.

The vertex separations of a graph G whose separators are finite form a subuniverse of $\vec{S}(V(G))$, the resulting submodular universe of all vertex separations of G together with the natural order function $(A, B) \mapsto |A, B| := |A \cap B|$, measuring the size of the *separator* $A \cap B$, is denoted by $\vec{S}(G)$.

A *tangle* in a graph G is an orientation τ of some $\vec{S}_k \subseteq \vec{S}(G)$ which has the *tangle property*:

$$\forall\ (A_1, B_1), (A_2, B_2), (A_3, B_3) \in P\colon G[A_1] \cup G[A_2] \cup G[A_3] \neq G. \tag{T}$$

Thus, the tangles in graph G are the $\mathcal{T}_G$-tangles in $\vec{S}(G)$ where

$$\mathcal{T}_G := \{\, \{(A_1, B_1), (A_2, B_2), (A_3, B_3)\} : (A_1, B_1), (A_2, B_2), (A_3, B_3) \in \vec{S}(G) \\ \text{with } G[A_1] \cup G[A_2] \cup G[A_3] = G \,\}.$$

Tangles in graphs are robust regular profiles in $\vec{S}(G)$. ([26])

For a graph G, order-respecting $S(G)$-trees and tree sets are an alternative view on another important object: a *tree-decomposition* $(T, \mathcal{V})$ of a graph G consists of a tree T and a family $\mathcal{V} = (V_t)_{t \in V(T)}$ of *bags* or *parts*, subsets of $V(G)$, so that

(T1) $V(G) = \bigcup_{t \in V(T)} V_t$;
(T2) for every edge $vw \in E(G)$ there is some $t \in V(T)$ so that $x, y \in V_t$;

(T3) for any $t_1, t_3 \in V(T)$ and t_2 on the unique path between t_1 and t_3 we have $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$.

The size of the largest bag minus one is the *width* of the tree-decomposition and the smallest width among all tree-decomposition of G is the *tree-width* of G.

The edges of a decomposition-tree T induce vertex separations of G: for an edge t_1t_2 of T let T_1 be the component of $T - t_1t_2$ which contains t_1 and let T_2 be the component which contains t_2, then $\left(\bigcup_{t\in V(T_1)} V_t, \bigcup_{t\in V(T_2)} V_t \right)$ is a vertex separation. The tree T of a tree-decomposition, together with its induced separations forms an order-respecting $S(G)$-tree. Vice versa, for a finite graphs G, every tree set in $\vec{S}(G)$ can also be turned into a tree-decomposition of G:

Lemma 2.3 ([26, 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:*

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

2.6.2. Bipartitions of a set

Given a set V, the subuniverse of $\vec{S}(V)$ consisting of only those set separations (A, B) where A and B are disjoint is called the separation system of *bipartitions of* V, we denote it by $\vec{\mathcal{B}}(V)$. Note that the only small bipartition is $(\emptyset, V)$ and no degenerate bipartitions exist unless V is empty.

For separations of this type, throughout the literature, often only one of the two sets, A or B, is denoted. In fact, a considerable branch of tangle theory is dedicated to tangles and branch-decompositions of *connectivity systems*, which have their origin in matroid theory [46]. A connectivity system is a submodular function $\lambda\colon 2^E \to \mathbb{Z}$ defined on the power set of some finite set E which is invariant under complementation, i.e., $\lambda(X) = \lambda(E \setminus X)$ for all $X \subseteq E$.

In the language of abstract separation systems, connectivity systems are (up to adding a large constant to the connectivity fuction to make it non-negative) universes of bipartitions equipped with a submodular order function. In this thesis we will follow the set-separation notation for bipartitions.

Part I.

Representing a single tangle indirectly

3. Tangles are decided by weighted vertex sets

In this chapter we show, that every tangle in a graph is determined by a weighting of the vertices. This is based on [36] and joint work with Christian Elbracht and Jakob Kneip.

The tangles in a graph G can be used to locate, and thereby capture the essence of, highly connected substructures in G in that every such substructure defines a tangle in G by orienting each low-order separation of G towards the side containing most or all of that substructure. If some tangle in G contains the separation (A, B), we think of A and B as the 'small' and the 'big' side of (A, B) in that tangle, respectively. The main result of this chapter will make this intuition concrete.

As a concrete example, if G contains an $n \times n$-grid for large n, then the vertex set of that grid defines a tangle τ in G as follows. Take note that no separation of low order can divide the grid into two parts of roughly equal size: If the grid is large enough then at least 90% of its vertices, say, will lie on the same side of such a separation. Orienting towards that side all the separations of order $< k$ for some fixed k much smaller than n then gives a tangle τ. In this way, the vertex set of the $n \times n$-grid 'determines τ by majority vote'.

In [26] Diestel raised the question whether all tangles in graphs arise in the above fashion, that is, whether all graph tangles are decided by majority vote by some subset of the vertices:

Problem 3.1. Given a k-tangle τ in a graph G, is there always a set X of vertices such that a separation (A, B) of order $< k$ lies in τ if and only if $|A \cap X| < |B \cap X|$?

A partial answer to this was given in [32], where Elbracht showed that such a set X always exists if G is $(k-1)$-connected and has at least $4(k-1)$ vertices. However, Elbracht's approach relies heavily on the $(k-1)$-connectedness of the graph and offers no line of attack for the general problem. Finding an answer for arbitrary graphs appears to be hard.

If a tangle in G is decided by some vertex set X by majority vote, this set X can be used as an oracle for that tangle, allowing one to store complete information about the complex structure of a tangle using a set of size at most $|V|$. On the other hand, if there were tangles without such a decider set, this would mean that tangles are a fundamentally more general concept than concrete highly cohesive subsets, not just an indirect way of capturing them.

In this chapter, we consider a fractional version of Diestel's question and answer it affirmatively, making precise the notion that B is the 'big' side of a separation

$(A, B) \in \tau$: given a k-tangle τ in G, rather than finding a vertex set X which decides τ by majority vote, we find a weight function $w\colon V(G) \to \mathbb{N}$ on the vertices such that for all separations (A, B) of order $< k$ we have $(A, B) \in \tau$ if and only if the vertices in B have higher total weight than those in A.

Thus, we show that every graph tangle is decided by some *weighted* set of vertices. This weight function, or weighted set of vertices, can then serve as an oracle for that tangle in the same way that a vertex set deciding the tangle by majority vote would. For any tangle, the existence of such a weight function with values in $\{0, 1\}$ is equivalent to the existence of a vertex set X deciding that tangle by majority vote.

Geelen [45] pointed out that the analogue of Diestel's question for tangles in matroids is false: there are matroid tangles which cannot be decided by majority vote, not even when considering a fractional version of the problem. Geelen's construction of such a matroid tangle is the matroid version of an example given in Section 3.2, where we show that another variant of tangles may fail to admit such a weight function as well.

In the next section we will formally define separations and tangles, and formulate and prove our main theorem. Following that, we show that the same arguments are also applicable to edge tangles of graphs, a relative of the tangles usually considered, and prove our main result also for this type of tangle.

3.1. Weighted deciders

Recall that a *separation* of a graph $G = (V, E)$ is a pair (A, B) with $A \cup B = V$ such that G contains no edge between $A \setminus B$ and $B \setminus A$, and the *order* of a separation (A, B) is the size $|A \cap B|$ of its separator $A \cap B$. Furthermore, for an integer k, a *k-tangle* in G is a set consisting of exactly one of (A, B) and (B, A) for every separation (A, B) of G of order $< k$, with the additional property that no three 'small' sides of separations in τ cover G, that is, that there are no $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$ for which $G = G[A_1] \cup G[A_2] \cup G[A_3]$.

Our main result is the following:

Theorem 1. *Let $G = (V, E)$ be a finite graph and τ a k-tangle in G. Then there exists a function $w\colon V \to \mathbb{N}$ such that a separation (A, B) of G of order $< k$ lies in τ if and only if $w(A) < w(B)$, where $w(U) := \sum_{u \in U} w(u)$ for $U \subseteq V$.*

We shall prove Theorem 1 in the remainder of this section. Our general strategy will be as follows: we define a partial order on the separations of G and consider the set of those separations of the k-tangle τ that are maximal in this partial order. For these separations we will be able to show that, on average, their separators divide each other so that they lie more on the 'big' side of each other, where 'big' is the big side according to τ. This will enable us to use a result from linear programming to find a weight function assigning weights to the vertices of these separators so that this weight function decides all these maximal separations of τ correctly. The nature of the partial order will then ensure that this weight function in fact decides all separations in τ correctly.

For a graph G there is a partial order on the separations of G given by letting $(A, B) \leqslant (C, D)$ if and only if $A \subseteq C$ and $B \supseteq D$. One of the main ingredients for the proof of Theorem 1 is the following observation about those separations in a tangle τ that are maximal in τ with respect to this partial order. It says, roughly, that they divide each other's separators so that, on average, those separators lie more on the big side of the separation than on the small side, according to the tangle.

Lemma 3.2. *For every k-tangle τ in a graph G and distinct maximal elements $(A, B), (C, D)$ of τ we have $|B \cap (C \cap D)| + |D \cap (A \cap B)| > |A \cap (C \cap D)| + |C \cap (A \cap B)|$.*

Proof. Let τ be a k-tangle in $G = (V, E)$, and let (A, B) and (C, D) be distinct maximal elements of τ. Observe that $(A \cup C\,,\, B \cap D)$ is a separation of G as well. In fact this separation is the supremum of (A, B) and (C, D) in the partial order. Therefore, τ cannot contain $(A \cup C\,,\, B \cap D)$ by the assumed maximality of (A, B) and (C, D) in τ. On the other hand τ cannot contain $(B \cap D\,,\, A \cup C)$ either since A, C, and $B \cap D$ together cover G. Consequently, since τ is a k-tangle, we must have $|(A \cup C) \cap (B \cap D)| \geqslant k$.

Recall that $|A \cap B| < k$ and $|C \cap D| < k$, since τ is a k-tangle. Observe additionally that the order of separations if modular, that is,

$$|A \cap B| + |C \cap D| = |(A \cup C) \cap (B \cap D)| + |(A \cap C) \cap (B \cup D)| \ .$$

The above inequalities thus imply that $|(A \cap C) \cap (B \cup D)| < k$, and hence in particular that $|(A \cap C) \cap (B \cup D)| < |(A \cup C) \cap (B \cap D)|$. Adding $|A \cap B \cap C \cap D|$ to both sides proves the claim. □

Additionally, we shall use a result from linear programming: Tucker's Theorem, a close relative of the Farkas Lemma. For a vector $x \in \mathbb{R}^n$ we use the usual shorthand notation $x \geqslant 0$ to indicate that all entries of x are non-negative, and similarly write $x > 0$ if all entries of x are strictly greater than zero.

Lemma 3.3 (Tucker's Theorem [67])**.** *Let $K \in \mathbb{R}^{n \times n}$ be a skew-symmetric matrix, i.e., $K^T = -K$. Then there exists a vector $x \in \mathbb{R}^n$ such that*

$$Kx \geqslant 0 \quad \textit{and} \quad x \geqslant 0 \quad \textit{and} \quad x + Kx > 0.$$

We are now ready to prove Theorem 1.

Theorem 1. Let a finite graph $G = (V, E)$ and a k-tangle τ in G be given. Since G is finite it suffices to find a weight function $w \colon V \to \mathbb{R}_{\geqslant 0}$ such that a separation

(A, B) of order $< k$ lies in τ precisely if $w(A) < w(B)$; by the density of the rationals in the reals, this w can then be turned into such a weight function with values in $\mathbb{N}$.

For this it is enough to find a function $w\colon V \to \mathbb{R}_{\geqslant 0}$ such that $w(A) < w(B)$ for all maximal elements (A, B) of τ: for if $w(A) < w(B)$ and $(C, D) \leqslant (A, B)$, then

$$w(C) \leqslant w(A) < w(B) \leqslant w(D).$$

So let us show that such a weight function w exists.

To this end let $(A_1, B_1), \dots, (A_n, B_n)$ be the maximal elements of τ and set

$$m_{ij} := \left|B_i \cap (A_j \cap B_j)\right| - \left|A_i \cap (A_j \cap B_j)\right|$$

for $i, j \leqslant n$. Let M be the matrix $(m_{ij})_{i,j \leqslant n}$. Observe that, by Lemma 3.2, we have $m_{ij} + m_{ji} > 0$ for all $i \neq j$, and hence the matrix $M + M^T$ has positive entries everywhere but on its diagonal (where it has zeros). We further define

$$K' := \frac{M + M^T}{2} \qquad \text{and} \qquad K := M - K'.$$

Then K is skew-symmetric, that is, $K^T = -K$. Let $x = (x_1, \dots, x_n)^T$ be the vector obtained by applying Lemma 3.3 to K. We define a weight function $w\colon V \to \mathbb{R}$ by

$$w(v) := \sum_{i\colon v \in A_i \cap B_i} x_i\,.$$

Note that w has its image in $\mathbb{R}_{\geqslant 0}$, and observe further that, for $Y \subseteq V$, we have

$$w(Y) = \sum_{y \in Y} w(y) = \sum_{i=1}^{n} x_i \cdot |Y \cap (A_i \cap B_i)|\,.$$

With this, for $i \leqslant n$, we have

$$\begin{aligned} w(B_i) - w(A_i) &= \sum_{j=1}^{n} x_j \cdot \left(\left|B_i \cap (A_j \cap B_j)\right| - \left|A_i \cap (A_j \cap B_j)\right|\right) \\ &= \sum_{j=1}^{n} x_j \cdot m_{ij} \\ &= (Mx)_i\,, \end{aligned}$$

where $(Mx)_i$ denotes the i-th coordinate of Mx. Thus w is the desired weight function if we can show that $Mx > 0$, that is, if all entries of Mx are positive.

From $x + Kx > 0$ we know that at least one entry of x is positive. Let us first consider the case that x has two or more positive entries. Then $K'x > 0$ since K' has positive values everywhere but on the diagonal, and hence

$$Mx = (K + K')\,x > 0$$

since $Kx \geqslant 0$. Therefore, in this case, w is the desired weight function.

Consider now the case that exactly one entry of x, say x_i, is positive, and that x is zero in all other coordinates. Then for $j \neq i$ we have $(Mx)_j \geqslant (K'x)_j > 0$ and thus $w(B_j) - w(A_j) = (Mx)_j > 0$. However $(Mx)_i = 0$ and thus $w(A_i) = w(B_i)$, so w is not yet as claimed. To finish the proof it remains to modify w so that $w(A_i) < w(B_i)$ while ensuring that we still have $w(A_j) < w(B_j)$ for $j \neq i$. This can be achieved by picking a sufficiently small $\varepsilon > 0$ so that $w(A_j) + \varepsilon < w(B_j)$ for all $j \neq i$, picking any $v \in B_i \setminus A_i$, and increasing the value of $w(v)$ by ε. □

We conclude with the remark that Theorem 1 and its proof extend to tangles in hypergraphs without any changes. Even more generally, the following more abstract version of Theorem 1 can be established with exactly the same proof:

Theorem 2. *Let U be a universe of set separations of a finite ground-set V with order function $|(A, B)| := |A \cap B|$. Then for any regular k-profile P in U there exists a function $w\colon V \to \mathbb{N}$ such that a separation (A, B) of order $< k$ lies in P if and only if $w(A) < w(B)$.*

Observe that if $G = (V, E)$ is a (hyper-)graph then the set U of all separations of G is such a universe of set separations. Moreover, every k-tangle τ of G is also a regular k-profile of U. (See [26] for more on the relation between graph tangles and profiles.) Therefore, Theorem 2 indeed applies to tangles in graphs and hypergraphs as well.

Theorem 2 holds with the same proof as Theorem 1, since Lemma 3.2 holds in this setting too: the only difference being that to see that $(B \cap D\,,\, A \cup C)$ cannot lie in the profile at hand one now has to use the definition of a regular k-profile rather than the fact that A, C, and $B \cap D$ cover G.

3.2. Edge tangles

A related object of study (cf. [27,61]) to the (vertex) tangles discussed above are the edge tangles of a graph. In this context one considers the *edge-cuts* of a (multi-)graph $G = (V, E)$, i.e., bipartitions (A, B) of V. The *order* of a cut (A, B) is the number of edges in G that are incident with vertices of both A and B. In this context one considers the *edge-cuts* of a (multi-)graph $G = (V, E)$, i.e., bipartitions (A, B) of V. The *order* of a cut (A, B) is the number of edges in G that are incident with vertices of both A and B. For an integer k, a *k-edge-tangle* of G is a set τ consisting of exactly one (A, B) or (B, A) for every cut (A, B) of order $< k$, with the additional properties that τ has no subset $\{\,(A_1, B_1), (A_2, B_2), (A_3, B_3)\,\}$ such that $B_1 \cap B_2 \cap B_3 = \emptyset$, and that τ contains no cut (A, B) for which B is incident with fewer than k edges of G.

In very much the same way as above we can prove the following theorem:

Theorem 3. *Let $G = (V, E)$ be a finite (multi-)graph and τ a k-edge-tangle in G. Then there exists a function $w\colon V \to \mathbb{N}$ such that a cut (A, B) of G of order $< k$ lies in τ if and only if $w(A) < w(B)$.*

We shall prove a more general version of this theorem where we allow G to be a graph with $\mathbb{R}_{\geqslant 0}$-weighted edges. We consider edges of weight 0 as indistinguishable from non-edges. Consequently, rather than a graph with weighted edges, we will just consider a pair (V, e) of a finite set V together with a symmetric function $e\colon V^2 \to \mathbb{R}_{\geqslant 0}$, which we shall call a *edge weighting.* The *order* of a bipartition (A, B) we define as $|(A, B)| := \sum_{(u,v)\in A\times B} e(u, v)$. Note that this function is submodular in the sense that

$$|(A, B)| + |(C, D)| \geqslant |(A \cup B, C \cap D)| + |(A \cap B, C \cup D)|.$$

For any positive r an *r-profile* in (V, e) is a set τ consisting of exactly one of (A, B) or (B, A) for every bipartition (A, B) of V of order $< r$, such that τ has no subset of the form $\{\,(A, B), (C, D), (B \cap D, A \cup C)\,\}$. We shall prove the following theorem, which directly implies Theorem 3:

Theorem 4. *Let (V, e) be an edge weighting and τ an r-profile in (V, e). Then there exists a function $w\colon V \to \mathbb{N}$ such that a bipartition (A, B) of V of order $< r$ lies in τ if and only if $w(A) < w(B)$.*

The main idea for proving this theorem is to first find an appropriate weighting of the edges by the same principles as in Theorem 1 and to then transform it into a weighted vertex decider. So let us first show an analogue of Lemma 3.2 for edge weightings. For this, we define a partial order on the bipartitions of V as in the previous section: by letting $(A, B) \leqslant (C, D)$ if and only if $A \subseteq C$ (and thus $B \supseteq D$). Using this partial order we can prove the following analogue of Lemma 3.2:

Lemma 3.4. *For every r-profile τ in a edge weighting (V, e) and distinct maximal elements $(A, B), (C, D)$ of τ we have*

$$\sum_{(u,v)\in B^2\cap(C\times D)} e(u, v) \;+ \sum_{(u,v)\in D^2\cap(A\times B)} e(u, v) > \sum_{(u,v)\in A^2\cap(C\times D)} e(u, v) \;+ \sum_{(u,v)\in C^2\cap(A\times B)} e(u, v)\,.$$

Proof. The bipartition $(A \cup C,\ B \cap D)$ of V is strictly larger in the partial order than the maximal elements (A, B) and (C, D) and hence cannot lie in τ. However, by the definition of an r-profile, τ cannot contain $(B \cap D,\ A \cup C)$ either. Thus, we must have $|(A \cup C,\ B \cap D)| \geqslant r$, from which it follows by submodularity that $|(A \cap C,\ B \cup D)| < r$. Combining these two inequalities, using the definition of order and adding $\sum_{u\in A\cap C} \sum_{v\in B\cap D} e(u, v)$ to both sides proves the claim. $\square$

We are now ready to prove Theorem 4:

Proof of Theorem 4. As in the proof of Theorem 1, it suffices to find a suitable weight function on pairs with values in $\mathbb{R}_{\geqslant 0}$. We will begin by finding a weight function $\overline{w}\colon V^2 \to \mathbb{R}_{\geqslant 0}$ such that we have $\overline{w}(A) \leqslant \overline{w}(B)$ for all $(A, B) \in \tau$ and $\overline{w}(A) < \overline{w}(B)$ for all but possibly one maximal element of τ, where $\overline{w}(A) = \sum_{(u,v)\in A^2} \overline{w}(u, v)$. Subsequently we use this to construct a suitable function $w\colon V \to \mathbb{R}_{\geqslant 0}$.

Enumerate the maximal elements of τ as $(A_1, B_1), \dots, (A_n, B_n)$. For every two maximal elements $(A_i, B_i), (A_j, B_j)$ let

$$m_{ij} := \sum_{(u,v)\in B_i^2 \,\cap\, A_j\times B_j} e(u, v) - \sum_{(u,v)\in A_i^2 \cap A_j\times B_j} e(u, v).$$

Let M be the matrix $(m_{ij})_{i,j\leqslant n}$. Observe that, by Lemma 3.4, $M + M^T$ has positive entries everywhere but on the diagonal, where it is zero. We are now in the same situation as in the proof of Theorem 1 and can find some vector x such that either $(Mx)_i > 0$ on all i, or x has exactly one non-zero entry, say x_i, and $(Mx)_j > 0$ for all $j \neq i$.

In either case, given a pair of vertices (u, v) let

$$\overline{w}(u, v) = e(u, v) \left(\sum_{j\colon (u,v)\in A_j\times B_j} x_j + \sum_{j\colon (u,v)\in B_j\times A_j} x_j \right) = \sum_{\substack{j \\ (u,v)\in(A_j\times B_j)\cup(B_j\times A_j)}} x_j\, e(u, v).$$

Note that $\overline{w}$ is symmetric. Define w as $w(v) = \sum_{u\in V} \overline{w}(v, u)$. We now have $w(B_i) - w(A_i) = 2(Mx)_i$, for indeed

$$\begin{aligned}
&w(B_i) - w(A_i) \\
&= \sum_{u\in B_i} \sum_{v\in V} \overline{w}(u, v) - \sum_{u\in A_i} \sum_{v\in V} \overline{w}(u, v) \\
&= \sum_{(u,v)\in B_i^2} \overline{w}(u, v) - \sum_{(u,v)\in A_i^2} \overline{w}(u, v) \\
&= \sum_{(u,v)\in B_i^2} \sum_{\substack{j \\ (u,v)\in(A_j\times B_j)\cup(B_j\times A_j)}} x_j\, e(u, v) \quad - \sum_{(u,v)\in A_i^2} \sum_{\substack{j \\ (u,v)\in(A_j\times B_j)\cup(B_j\times A_j)}} x_j\, e(u, v) \\
&= 2\sum_j \left(\sum_{(u,v)\in B_i^2\cap(A_j\times B_j)} x_j\, e(u, v) \;-\; \sum_{(u,v)\in A_i^2\cap(A_j\times B_j)} x_j\, e(u, v) \right) \\
&= 2(Mx)_i.
\end{aligned}$$

Thus either $w(B_i) > w(A_i)$ for all maximal elements of τ, from which the claim follows directly, or there is a single maximal element (A_i, B_i) of τ such that $w(B_i) = w(A_i)$ and $w(B_j) > w(A_j)$ for all others. However, as in the proof of Theorem 1, in the latter case we can pick an arbitrary vertex $v \in B_i$ and increase $w(v)$ by some small $\varepsilon > 0$ to achieve $w(B_i) > w(A_i)$ while keeping $w(B_j) > w(A_j)$ for all other maximal elements of τ. $\square$

Remarkably, and in contrast to Theorem 1, Theorem 3 does not in fact extend to hypergraphs. To demonstrate this let us recall the relevant definitions.

A hypergraph $H = (V, E)$ consists of a vertex set V together with a set $E \subseteq 2^V$ of hyperedges. An *(edge) cut* of a hypergraph $H = (V, E)$ is a bipartition (A, B) of V. The *order* of such an edge cut (A, B) is the number of hyperedges of H that are incident with vertices from both A and B.

For an integer k, a *k-edge-tangle* of H is a set τ consisting of exactly one (A, B) or (B, A) for every cut (A, B) of order $< k$, with the additional properties that τ has no subset $\{ (A_1, B_1), (A_2, B_2), (A_3, B_3) \}$ such that $B_1 \cap B_2 \cap B_3 = \emptyset$, and that τ contains no cut (A, B) for which B is incident with fewer than k hyperedges of H.

A *weighted decider* for some k-edge-tangle τ of a hypergraph $H = (V, E)$ is a function $w\colon V \to \mathbb{N}$ such that a cut (A, B) of H of order $< k$ lies in τ if and only if $w(A) < w(B)$, where we use the shorthand notation $w(U) := \sum_{u \in U} w(u)$ for $U \subseteq V$.

Theorem 3 then asserts that if H is a simple graph, i.e., if every hyperedge in E has size 2, then every k-edge-tangle of H has such a weighted decider. We are now going to construct an example demonstrating that this may fail for hypergraphs H that are not simple graphs.

Example 3.5. For some natural number $k \geqslant 6$ let ℓ be an integer with $3 \leqslant \ell \leqslant \frac{k}{2}$. Let V be the set of all ℓ-element subsets of $[k] = \{ 1, \dots, k \}$. Let the set E of hyperedges consist of, for each $i \in [k]$, the set of all $v \in V$ that contain i.

Note that each of these k many hyperedges of H has size $\binom{k-1}{\ell-1}$, making H a uniform ℓ-regular hypergraph.

We now show that each hypergraph of this form has a tangle without a decider set.

Theorem 3.6. *Let H be as in Example 3.5. Then H has a k-edge-tangle with no weighted decider.*

Proof. Let S_k denote the set of all cuts of H of order $< k$. For a set $A \subseteq V$ we note that $\bigcup A$ is the set $\bigcup_{v \in A} v$, which is a subset of $[k]$. Observe that for every cut (A, B) of H at most one of $\bigcup A$ and $\bigcup B$ can be a proper subset of $[k]$. Note further that a cut (A, B) of H lies in S_k if and only if at least one of the k hyperedges of H does not meet both A and B, which is the case precisely if one of $\bigcup A$ and $\bigcup B$ is a proper subset of $[k]$.

We can therefore define

$$\tau := \{ (A, B) \in S_k : \bigcup A \subsetneq [k] \} .$$

Let us show that τ is a k-edge-tangle of H with no weighted decider.

To see that τ is a k-edge-tangle we note that by the above observation τ contains exactly one of (A, B) or (B, A) for every cut $(A, B) \in S_k$. Furthermore if

$(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$ then any element of V containing at least one point each from $[k] \setminus \bigcup A_1$, from $[k] \setminus \bigcup A_2$, and from $[k] \setminus \bigcup A_3$ lies in $B_1 \cap B_2 \cap B_3$, which is hence non-empty since such a $v \in V$ exists by $\ell \geqslant 3$. Finally for each $(A, B) \in \tau$ the set B is incident with each hyperedge of H since $\bigcup B = [k]$. Thus τ is indeed a k-edge-tangle.

Finally, let us show that τ has no weighted decider. Suppose for a contradiction that some weighted decider $w\colon V \to \mathbb{N}$ for τ exists. For each $i \in [k]$ consider the cut (A_i, B_i), where

$$A_i := \{\, v \in V : i \notin v \,\} \qquad \text{and} \qquad B_i := \{\, v \in V : i \in v \,\},$$

and note that $(A_i, B_i) \in \tau$. Since w is a weighted decider for τ we have $w(B_i) > w(A_i)$ for each $i \in [k]$. We therefore have

$$\sum_{i \in [k]} (w(B_i) - w(A_i)) > 0\,,$$

since each term in the sum is positive. By counting the instances of $w(v)$ occurring in the sum for each $v \in V$ we find that

$$\sum_{i \in [k]} (w(B_i) - w(A_i)) = \sum_{v \in V} w(v) \cdot \left(|\{\, i \in [k] : i \in v \,\}| - |\{\, i \in [k] : i \notin v \,\}| \right)\,,$$

since $v \in B_i$ if and only if $i \in v$, and otherwise $v \in A_i$. The left-hand side of this equation is positive. However, in contradiction to this, no term of the right-hand sum is greater than zero since we have by $\ell \leqslant \frac{k}{2}$ that

$$|\{\, i \in [k] : i \in v \,\}| - |\{\, i \in [k] : i \notin v \,\}| = \ell - (k - \ell) \leqslant 0\,.$$

Therefore there can be no weighted decider for τ. $\square$

A construction analogous to Example 3.5 was found independently by Geelen [45] in the setting of matroids, who used it to show that matroids, too, can have tangles with no weighted decider.

Finally, let us remark that Example 3.5 can also be used to show that allowing weighted deciders to take values in $\mathbb{R}$ rather than $\mathbb{N}$ does not suffice to guarantee their existence for edge tangles of hypergraphs: for $k = 2\ell$ the tangle described in Theorem 3.6 has no weighted decider with real-valued and possibly negative weights either, with the same proof.

4. Dual tangles on a bipartite graph

In this chapter we present the dual separation systems and tangles that naturally occur on a bipartite graph. This is based on part of [24] which is still in the process of writing in collaboration with Christian Elbracht, Reinhard Diestel, and Joshua Erde.

Many relations in our everyday lives can naturally be expressed in the form of a bipartite graph. Consider, for example, in an online shop the containment-relation between products and purchases, or for a collection of recipes the 'is an ingredient'-relation between recipes and items of produce. Given such a context, we usually understand how 'cohesive' a subset of one class, X, is in terms of elements of the other class, Y, which bear a common relation to them. For example, flour, milk, eggs, sugar, and baking powder obviously form a coherent group of ingredients, as witnessed by them belonging to the recipes for various kinds of cake. The recipes for cakes, in turn, form a coherent group of recipes as witnessed by those common ingredients.

In this chapter we will present a way to formalize tangles on the sides of any given bipartite graph which captures this concept of duality. It will turn out that the dual nature of X and Y will allow us to give quite a natural order function on arbitrary set separations of X or Y, which we will prove is submodular. It is then natural to study the tangle structure of the set of separations of X or Y of order less than k for fixed k. We will show that a – quite natural – correspondence between separations of X and Y will extend to a correspondence between the low order tangles of the separation of X and Y, allowing us to relate the tangle structure of X with that of Y.

4.1. Tangles on the sides of a bipartite graph

To make what we discussed above precise, let $G = (V, E)$ be a fixed bipartite connected graph with partition classes X and Y; we will keep these fixed for the entire chapter. Recall, that we denote by $S(X)$ the set of all set separations of X, that is the set of all sets $\{ A, B \}$ with $A, B \subseteq X$ such that $A \cup B = X$. Similarly, we denote by $S(Y)$ the set of all set separations of Y, and we denote by $\vec{S}(X)$ and $\vec{S}(Y)$ the set of oriented separations from $S(X)$ or $S(Y)$ respectively.

Then, the structure of the bipartite graph G allow us to relate the separations in $\vec{S}(X)$ to the separations in $\vec{S}(Y)$, i.e., we will obtain a *dual separation* to a given separation in $\vec{S}(X)$. One natural way to do so is as follows: given a separation

(A, B) of X there will be some vertices in Y which are joined in G to more vertices in A than in B, while other vertices in Y are joined to more vertices in B than in A. This defines a way of partitioning the vertices in Y. So, given $(A, B) \in \vec{S}(X)$ we define the separation $(A, B)^{\triangleright} := (A_B^{\triangleright}, B_A^{\triangleright}) \in \vec{S}(Y)$ by letting

$$A_B^{\triangleright} := \{\, y \in Y : |N(y) \cap A| \geqslant |N(y) \cap B| \,\}$$

and

$$B_A^{\triangleright} := \{\, y \in Y : |N(y) \cap A| \leqslant |N(y) \cap B| \,\}.$$

We call $(A, B)^{\triangleright}$ the *shift* of (A, B).

Similarly,[1] a set separation (C, D) of Y gives rise to its *shift* $(C, D)^{\triangleleft} := (C_D^{\triangleleft}, D_C^{\triangleleft})$, a set separation of X, via

$$C_D^{\triangleleft} := \{\, x \in X : |N(x) \cap C| \geqslant |N(x) \cap D| \,\}$$

and

$$D_C^{\triangleleft} := \{\, x \in X : |N(x) \cap C| \leqslant |N(x) \cap D| \,\}.$$

We note that both these shifting operations commute with the natural involutions on $\vec{S}(X)$ and $\vec{S}(Y)$. However, we also note that this operation is not necessarily idempotent: there may exist $(A, B) \in \vec{S}(X)$ such that $((A, B)^{\triangleright})^{\triangleleft} \neq (A, B)$.

The map $(\cdot)^{\triangleright} \colon \vec{S}(X) \to \vec{S}(Y)$ induces an inverse 'pull-back' map ${}^{\triangleleft}(\cdot) \colon 2^{\vec{S}(Y)} \to 2^{\vec{S}(X)}$, sending every $\tau \subseteq \vec{S}(Y)$ to

$${}^{\triangleleft}\tau := \{\, (A, B) \in \vec{S}(X) \colon (A, B)^{\triangleright} \in \tau \,\} \subseteq \vec{S}(X).$$

Similarly, the map $(\cdot)^{\triangleleft} \colon \vec{S}(Y) \to \vec{S}(X)$ induces a map ${}^{\triangleright}(\cdot) \colon 2^{\vec{S}(X)} \to 2^{\vec{S}(Y)}$ sending every $\tau \subseteq \vec{S}(X)$ to

$${}^{\triangleright}\tau := \{\, (C, D) \in \vec{S}(Y) \colon (C, D)^{\triangleleft} \in \tau \,\} \subseteq \vec{S}(Y).$$

The question then arises, under which conditions on a tangle τ will the subset ${}^{\triangleright}\tau$ or ${}^{\triangleleft}\tau$ also be a tangle? In order for there to be any interesting tangle structure we will have to restrict to some subset of $\vec{S}(X)$ or $\vec{S}(Y)$, and the most natural way to do so will be to choose some order function and consider the separations $\vec{S}_k(X)$ or $\vec{S}_k(Y)$ of order less than k. However, then in order for the pull-back to have any hope of being a tangle, it must orient every separation in $\vec{S}_{k'}(X)$ or $\vec{S}_{k'}(Y)$ for some k'. Hence, already for this question to make sense, we will need to choose

[1] Informally, we think of the vertex classes X, Y of G as being its 'left' and 'right' class, respectively. Formally, however, $\{\, X, Y \,\}$ is an unordered pair, so the operators $(\cdot)^{\triangleright}$ and $(\cdot)^{\triangleleft}$ are formally the same: they map their argument, an oriented separation of one of the sets X, Y, to an oriented separation of the other set. It is important, that we never treat X and Y differently: they are disjoint, but indistinguishable.

an appropriate order function which behaves nicely with respect to the shifting operation.

In fact, we will define order functions on $\vec{S}(X)$ and $\vec{S}(Y)$ so that shifting a separation never increases its order. This will guarantee that if τ orients all the separations of order less than k in $\vec{S}(X)$, then ${}^{\triangleright}\tau$ orients all the separations of order less than k in $\vec{S}(Y)$. Indeed, if $(C,D) \in \vec{S}(Y)$ has order less than k, then $(A,B) := (C,D)^{\triangleleft} \in \vec{S}(X)$ has order less than k and so precisely one of (A,B) or (B,A) is in τ by assumption. Since $(B,A) = (D,C)^{\triangleleft}$ it follows that precisely one of (C,D) or (D,C) is in ${}^{\triangleright}\tau$.

Furthermore, these order functions are defined in a particularly natural way, determined only by the structure of G. Broadly, the order functions measure in some way how evenly a separation of X or Y splits the neighbourhood of each vertex from the appropriate class. For example, in our online shop example, the order of a separation (A,B) of V will be determined by how evenly this separation splits the set of items bought in each purchase. The more balanced the split, the larger the contribution of this vertex to the order of the separation. In this way, separations for which most vertices in the opposite partition class have a clear 'preference' of one side or the other will have low order.

Explicitly, let us define the order function $|\cdot|_X \colon \vec{S}(X) \to \frac{1}{2}\mathbb{N}$ where

$$|A,B|_X := \sum_{y\in Y} \left(\mathsf{min}\{\, |N(y)\cap A|\,, |N(y)\cap B|\,\} - \frac{|N(y)\cap A\cap B|}{2} \right).$$

Here the first term will be larger when $N(y)$ is more evenly split by (A,B). Similarly, we define $|\cdot|_Y \colon \vec{S}(Y) \to \frac{1}{2}\mathbb{N}$ where

$$|C,D|_Y := \sum_{x\in X} \left(\mathsf{min}\{\, |N(x)\cap C|\,, |N(x)\cap D|\,\} - \frac{|N(x)\cap C\cap D|}{2} \right).$$

Note that these functions are symmetric and non-negative, as required of an order function for separation systems. Moreover, the function $|\cdot|_X$ attains its maximum value on the separation (X,X), and since orientations of all of $\vec{S}(X)$ are not enlightening, we will in the following assume implicitly that any $\vec{S}_k(X)$ we consider does not contain the separation (X,X).

Less obviously, these order functions are submodular, so $\vec{S}(X)$ and $\vec{S}(Y)$ equipped with these functions are submodular universes. Submodularity is a fundamental property for order functions at the heart of tangle theory, and so we include the proof even though it is straightforward. However, the reader is invited to skip the proofs of the next two lemmas at first reading, to remain with the flow of the narrative.

Proposition 4.1. *The order function $|\cdot|_X$ is submodular.*

Proof. We show that $|\cdot|_X$ is a sum of submodular functions. For this consider, for $y \in Y$, the order function on $\vec{S}(X)$ given by

$$|A,B|_y := \mathsf{min}\{\,|N(y)\cap A|\,,|N(y)\cap B|\,\} - \frac{|N(y)\cap A\cap B|}{2}$$

and note that $|A,B|_X = \sum_{y\in Y}|A,B|_y$, thus it is enough to show that $|\cdot|_y$ is submodular for every $y \in Y$. Fix some y in Y. For $Z \subseteq X$ we denote $N_Z := |N(y)\cap Z|$.

Let (A_1,B_1) and (A_2,B_2) be separations in $\vec{S}(X)$, and suppose without loss of generality that $N_{A_i} \leqslant N_{B_i}$. Let $A_i' := A_i \setminus B_i$, $B_i' := B_i \setminus A_i$ and $Z_i := A_i \cap B_i$. Note that $|A_i,B_i|_y = N_{A_i'} + N_{Z_i}/2$.

We observe that

$$|A_1\cap A_2, B_1\cup B_2|_y = N_{A_1'\cap A_2'} + \frac{1}{2}(N_{Z_1\cap Z_2} + N_{Z_1\cap A_2'} + N_{A_1'\cap Z_2}),$$

and

$$|A_1\cup A_2, B_1\cap B_2|_y = \mathsf{min}\{\,N_{A_1'\cup A_2'}, N_{B_1'\cap B_2'}\,\} + \frac{1}{2}(N_{Z_1\cap Z_2} + N_{Z_1\cap B_2'} + N_{B_1'\cap Z_2}).$$

Summing these two, we get

$$\begin{aligned}
&|A_1\cap A_2, B_1\cup B_2|_y + |A_1\cup A_2, B_1\cap B_2|_y\\
&= N_{A_1'\cap A_2'} + \frac{1}{2}(N_{Z_1\cap Z_2} + N_{Z_1\cap A_2'} + N_{A_1'\cap Z_2})\\
&\qquad + \mathsf{min}\{\,N_{A_1'\cup A_2'}, N_{B_1'\cap B_2'}\,\} + \frac{1}{2}(N_{Z_1\cap Z_2} + N_{Z_1\cap B_2'} + N_{B_1'\cap Z_2})\\
&\leqslant N_{A_1'\cap A_2'} + N_{A_1'\cup A_2'}\\
&\qquad + \frac{1}{2}(N_{Z_1\cap Z_2} + N_{Z_1\cap A_2'} + N_{Z_1\cap B_2'} + N_{Z_1\cap Z_2} + N_{A_1'\cap Z_2} + N_{B_1'\cap Z_2})\\
&= N_{A_1'} + N_{A_2'} + \frac{1}{2}(N_{Z_1} + N_{Z_2})\\
&= |A_1,B_1|_y + |A_2,B_2|_y\,.
\end{aligned}$$

Similarly,

$$|A_1\cap B_2, B_1\cup A_2|_y = \mathsf{min}\{\,N_{A_1'\cap B_2'}, N_{B_1'\cup A_2'}\,\} + \frac{1}{2}(N_{Z_1\cap Z_2} + N_{Z_1\cap B_2'} + N_{A_1'\cap Z_2}),$$

and

$$|A_1\cup B_2, B_1\cap A_2|_y = \mathsf{min}\{\,N_{A_1'\cup B_2'}, N_{B_1'\cap A_2'}\,\} + \frac{1}{2}(N_{Z_1\cap Z_2} + N_{Z_1\cap A_2'} + N_{B_1'\cap Z_2}).$$

Summing these two, we get

$$\begin{aligned}
&|A_1\cap B_2, B_1\cup A_2|_y + |A_1\cup B_2, B_1\cap A_2|_y\\
&= \mathsf{min}\{\,N_{A_1'\cap B_2'}, N_{B_1'\cup A_2'}\,\} + \frac{1}{2}(N_{Z_1\cap Z_2} + N_{Z_1\cap B_2'} + N_{A_1'\cap Z_2})\\
&\qquad + \mathsf{min}\{\,N_{A_1'\cup B_2'}, N_{B_1'\cap A_2'}\,\} + \frac{1}{2}(N_{Z_1\cap Z_2} + N_{Z_1\cap A_2'} + N_{B_1'\cap Z_2})
\end{aligned}$$

$$\begin{aligned}
&\leqslant N_{A_1'\cap B_2'} + N_{A_2'\cap B_1'} \\
&\qquad + \frac{1}{2}(N_{Z_1\cap Z_2} + N_{Z_1\cap A_2'} + N_{Z_1\cap B_2'} + N_{Z_1\cap Z_2} + N_{A_1'\cap Z_2} + N_{B_1'\cap Z_2}) \\
&\leqslant N_{A_1'} + N_{A_2'} + \frac{1}{2}(N_{Z_1} + N_{Z_2}) \\
&= |A_1, B_1|_y + |A_2, B_2|_y\,.
\end{aligned}$$

Thus $|\cdot|_y$ is submodular, and so is $|\cdot|_X = \sum_{y\in Y} |\cdot|_y$. $\square$

Next, we show that the shifting operation does not increase the order of a separation. For this we first show the following lemma, giving an alternative representation of the order function:

Lemma 4.2. *For all $(A,B) \in \vec{S}(X)$ we have*

$$|A,B|_X = \left|E(A_B^{\triangleright}, B)\right| + \left|E(B_A^{\triangleright}, A)\right| - \left|E(A_B^{\triangleright} \cap B_A^{\triangleright}, X)\right|/2 - |E(Y, A\cap B)|/2.$$

Proof. This can be calculated by rearranging sums:

$$\begin{aligned}
&|A,B|_X \\
&= \sum_{y\in Y}\left(\min\{\,|N(y)\cap A|\,,|N(y)\cap B|\,\} - \frac{|N(y)\cap A\cap B|}{2}\right) \\
&= \sum_{\substack{y\in Y\\ |N(y)\cap A|\geqslant|N(y)\cap B|}} |N(y)\cap B| + \sum_{\substack{y\in Y\\ |N(y)\cap B|\geqslant|N(y)\cap A|}} |N(y)\cap A| \\
&\qquad -\Bigg(\sum_{\substack{y\in Y\\ |N(y)\cap A|=|N(y)\cap B|}} \frac{|N(y)|}{2} + \sum_{y\in Y}\frac{|N(y)\cap A\cap B|}{2}\Bigg) \\
&= \left|E(A_B^{\triangleright}, B)\right| + \left|E(B_A^{\triangleright}, A)\right| - \left|E(A_B^{\triangleright} \cap B_A^{\triangleright}, X)\right|/2 - |E(Y, A\cap B)|/2. \qquad \square
\end{aligned}$$

With this we can now prove that shifting a separation indeed cannot increase the order of a separation:

Lemma 4.3. *Let (A,B) be a separation of X, then $|A,B|_X \geqslant |A_B^{\triangleright}, B_A^{\triangleright}|_Y$. Similarly if (C,D) is a separation of Y, then $|C,D|_Y \geqslant |C_D^{\triangleleft}, D_C^{\triangleleft}|_X$.*

Proof. This is true by the following calculation:

$$\begin{aligned}
&|A_B^{\triangleright}, B_A^{\triangleright}|_Y \\
&= \sum_{x\in X}\left(\min\{\left|N(x)\cap A_B^{\triangleright}\right|, \left|N(x)\cap B_A^{\triangleright}\right|\} - \frac{\left|N(x)\cap A_B^{\triangleright}\cap B_A^{\triangleright}\right|}{2}\right) \\
&\leqslant \sum_{a\in A}\left|N(a)\cap B_A^{\triangleright}\right| + \sum_{b\in B}\left|N(b)\cap A_B^{\triangleright}\right| - \sum_{x\in A\cap B}\frac{|N(x)|}{2} - \sum_{y\in A_B^{\triangleright}\cap B_A^{\triangleright}}\frac{|N(y)|}{2}
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{b \in B_A^{\triangleright}} |N(b) \cap A| + \sum_{a \in A_B^{\triangleright}} |N(a) \cap B| - \sum_{y \in A_B^{\triangleright} \cap B_A^{\triangleright}} \frac{|N(y)|}{2} - \sum_{x \in A \cap B} \frac{|N(x)|}{2} \\
&= E(A_B^{\triangleright}, B) + E(B_A^{\triangleright}, A) - E(A_B^{\triangleright} \cap B_A^{\triangleright}, X)/2 - E(Y, A \cap B)/2 \\
&= |A, B|_X. \qquad \square
\end{aligned}
$$

Finally, in order to define the tangles of $\vec{S}(X)$ and $\vec{S}(Y)$ we need to define the notion of consistency[2] that we require our orientations to satisfy. There are a few natural choices that one could make here, however in most contexts it turns out that these definitions are in some sense weakly equivalent, in that tangles under any one definition tend to induce tangles of slightly lower order under the other definitions.

With that in mind, let us define a *tangle of* $\vec{S}_k(X)$ *(in* G*)* as an orientation τ of $\vec{S}_k(X)$ which satisfies the following property:

$$\textit{There are no } (A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau \textit{ with } A_1 \cup A_2 \cup A_3 = X. \tag{†}$$

We define tangles of $\vec{S}_k(Y)$ in G accordingly. This is perhaps the simplest definition to take, and is a direct analogue of the corresponding notion of 'consistency' used to define tangles in matroids. We will discuss later in more detail the extent to which our results hold for tangles defined in terms of other notions of 'consistency'.

Let us introduce one more piece of shorthand notation:

$$^{\triangleright}_{k}\tau := {}^{\triangleright}\tau \cap \vec{S}_k(Y) \quad \text{ and } \quad ^{\triangleleft}_{k}\tau := {}^{\triangleleft}\tau \cap \vec{S}_k(X).$$

With that, we are ready to state the main results of this section, which, with the aid of the order function and shifting operation, relate the tangles of $\vec{S}(X)$ to those of $\vec{S}(Y)$.

Theorem 5. *Let τ be a tangle of $\vec{S}_{4k}(X)$, then $\tau' := {}^{\triangleright}_{k}\tau$ is a tangle of $\vec{S}_k(Y)$.*

Proof. We first note that τ' is an orientation of $\vec{S}_k(Y)$. Indeed, suppose that both (C, D) and (D, C) are in τ'. If we let

$$A = \{\, x \in X \colon\ |N(x) \cap C| \geqslant |N(x) \cap D| \,\}$$

and

$$B = \{\, x \in X \colon\ |N(x) \cap C| \leqslant |N(x) \cap D| \,\},$$

then $(C, D)^{\triangleleft} = (A, B)$ and $(D, C)^{\triangleleft} = (B, A)$ and by assumption both of these separations are in τ contradicting the fact that τ is an orientation.

So, it remains to show that τ' satisfies (†). Let us suppose for contradiction that there is some set $\{\,(C_1, D_1), (C_2, D_2), (C_3, D_3)\,\} \subseteq \tau'$ such that $C_1 \cup C_2 \cup C_3 = Y$.

[2]'Consistency' here is meant not in the technical sense of no two separations pointing away from each other, but in the broader sense of 'a tangle-like property'.

Let $(A_i, B_i) = (C_i, D_i)^{\triangleleft}$ for each $i = 1, 2, 3$. Then, since $(A_i, B_i) \in \tau$ for each i, and τ is a tangle, it follows that there is some non-empty set Z such that $Z = X \setminus (A_1 \cup A_2 \cup A_3)$.

Since $Z \subseteq B_i$ for each i, $|N(x) \cap D_i| \geqslant |N(x) \cap C_i|$ for all $x \in Z$ and $i = 1, 2, 3$. However, since $C_1 \cup C_2 \cup C_3 = Y$,

$$\begin{aligned}
\sum_{i=1}^{3} |C_i, D_i|_Y &= \sum_{i=1}^{3} \sum_{x \in X} \left(\mathsf{min}\{\, |N(x) \cap C_i| \,, |N(x) \cap D_i| \,\} - \frac{|N(x) \cap C_i \cap D_i|}{2} \right) \\
&\geqslant \sum_{i=1}^{3} \sum_{x \in Z} \left(\mathsf{min}\{\, |N(x) \cap C_i| \,, |N(x) \cap D_i| \,\} - \frac{|N(x) \cap C_i \cap D_i|}{2} \right) \\
&\geqslant \sum_{x \in Z} \sum_{i=1}^{3} \left(|N(x) \cap C_i| - \frac{|N(x) \cap C_i \cap D_i|}{2} \right) \\
&\geqslant \sum_{z \in Z} \frac{d(z)}{2} = \frac{|E(Z, Y)|}{2}.
\end{aligned}$$

Hence $|E(Z, Y)| < 6k$. Then,

$$|Z, X|_X = \frac{E(Z, Y)}{2} \leqslant 3k.$$

Hence $(Z, X) \in \tau$ by (†).

Finally, since $|A_3, B_3|_X \leqslant |C_3, D_3|_Y < k$ and $|E(Z, Y)| \leqslant 6k$ we can conclude by Proposition 4.1 that

$$|A_3 \cup Z, B_3 \cap X|_X \leqslant |A_3, B_3|_X + |Z, X|_X < k + 3k = 4k.$$

Hence, since $(A_3, B_3), (Z, X) \in \tau$ and $|A_3 \cup Z, B_3|_X < 4k$, it follows from (†) that $(A_3 \cup Z, B_3) \in \tau$. However, then

$$\{\, (A_1, B_1), (A_2, B_2), (A_3 \cup Z, B_3) \,\} \subseteq \tau$$

and $A_1 \cup A_2 \cup (A_3 \cup Z) = X$, contradicting (†). $\square$

By symmetry, we then obtain a similar conclusion as in Theorem 5 when we shift a tangle of $\vec{S}_{4k}(Y)$.

A natural question then to ask at this point, is, even if the shifting operations themselves are not idempotent, whether the operation they induce on tangles is in some way 'idempotent': if we shift a tangle twice, do we end up with the original tangle? It turns out that, again up to a constant factor, this is indeed the case.

Theorem 6. *Let τ be a tangle of $\vec{S}_{16k}(X)$, let $\tau' = {}^{\triangleright}_{4k}\tau$, and let $\tau'' = {}^{\triangleleft}_{k}\tau'$. Then $\tau'' \subseteq \tau$.*

To prove this theorem, we first need to analyse how a separation of X can behave under shifting that separation from X to Y and then back to X. It turns out that behaviour of this 'double shift' depends on the relation between the order of the separation and its shift. Our first lemma analyses the case that these two orders are the same:

Lemma 4.4. *If $|A,B|_X = |(A,B)^{\triangleright}|_Y$ then $A \subseteq (A^{\triangleright}_B)^{\triangleleft}_{B^{\triangleright}_A}$ and $B \subseteq (B^{\triangleright}_A)^{\triangleleft}_{A^{\triangleright}_B}$.*

Proof. By the proof of Lemma 4.3, we have that

$$\begin{aligned}
&|A^{\triangleright}_B, B^{\triangleright}_A|_Y \\
&= \sum_{x\in X}\left(\min\{\,\left|N(x)\cap A^{\triangleright}_B\right|, \left|N(x)\cap B^{\triangleright}_A\right|\,\} - \frac{\left|N(x)\cap A^{\triangleright}_B\cap B^{\triangleright}_A\right|}{2}\right)\\
&\leqslant \sum_{a\in A}\left|N(a)\cap B^{\triangleright}_A\right| + \sum_{b\in B}\left|N(b)\cap A^{\triangleright}_B\right| - \sum_{x\in A\cap B}\frac{|N(x)|}{2} - \sum_{y\in A^{\triangleright}_B\cap B^{\triangleright}_A}\frac{|N(y)|}{2}\\
&= |A,B|_X.
\end{aligned}$$

Thus, if $|A,B|_X = |A^{\triangleright}_B, B^{\triangleright}_A|_Y$, then

$$\begin{aligned}
&\sum_{x\in X}\min\{\,\left|N(x)\cap A^{\triangleright}_B\right|, \left|N(x)\cap B^{\triangleright}_A\right|\,\}\\
&= \sum_{a\in A}\left|N(a)\cap B^{\triangleright}_A\right| + \sum_{b\in B}\left|N(b)\cap A^{\triangleright}_B\right| - \sum_{x\in A\cap B}\frac{|N(x)|}{2}.
\end{aligned}$$

In particular, for $x \in A$ we have $\left|N(x)\cap B^{\triangleright}_A\right| \leqslant \left|N(x)\cap A^{\triangleright}_B\right|$ and thus $x \in (A^{\triangleright}_B)^{\triangleleft}_{B^{\triangleright}_A}$. Similarly, for $x \in B$ we have $\left|N(x)\cap B^{\triangleright}_A\right| \geqslant \left|N(x)\cap A^{\triangleright}_B\right|$ and thus $x \in (B^{\triangleright}_A)^{\triangleleft}_{A^{\triangleright}_B}$. □

While the previous lemma analysed the case that the order of the shift equals the order of the separation we started with, the next two lemmas allows us to obtain additional information when this is not the case.

Lemma 4.5. *For every $x \in A\setminus B$ with $\left|N(x)\cap B^{\triangleright}_A\right| > \left|N(x)\cap A^{\triangleright}_B\right|$ (equivalently: for every $x \in (B^{\triangleright}_A)^{\triangleleft}_{A^{\triangleright}_B} \setminus B$) we have $|A+x,\, B+x|_X \leqslant |A,B|_X$ and $|\{x\},X|_X \leqslant |A,B|_X$.*

Symmetrically, the same is true for every $x \in B \setminus A$ where $\left|N(x)\cap A^{\triangleright}_B\right| > \left|N(x)\cap B^{\triangleright}_A\right|$, or equivalently, every $x \in (A^{\triangleright}_B)^{\triangleleft}_{B^{\triangleright}_A} \setminus A$.

Proof. We have

$$\begin{aligned}
&|A, B+x|_X \\
&= \sum_{y\in Y}\left(\mathsf{min}\{\,|N(y)\cap A|\,, |N(y)\cap (B\cup\{x\})|\,\} - \frac{|N(y)\cap A\cap (B+x)|}{2}\right) \\
&= \sum_{y\in Y}\left(\mathsf{min}\{\,|N(y)\cap A|\,, |N(y)\cap B|\,\} - \frac{|N(y)\cap A\cap B|}{2}\right) \\
&\qquad + E(A_B^{\triangleright}\setminus B_A^{\triangleright}, \{x\}) - \frac{|N(x)|}{2} \\
&< |A, B|_X.
\end{aligned}$$

Moreover $|\{x\}, X|_X = \frac{|N(x)|}{2}$ and for every $y \in N(x)\cap B_A^{\triangleright}$ we have that

$$\begin{aligned}
\mathsf{min}\{\,|N(y)\cap A|\,, |N(y)\cap B|\,\} &= |N(y)\cap A| \\
&\geqslant 1 + |N(y)\cap (A-x)| \geqslant 1 + |N(y)\cap A\cap B|\,,
\end{aligned}$$

thus $\mathsf{min}\{\,|N(y)\cap A|\,, |N(y)\cap B|\,\} - \frac{|N(y)\cap A\cap B|}{2} \geqslant 1$, which gives

$$\begin{aligned}
|A, B|_X &\geqslant \sum_{y\in N(x)\cap B_A^{\triangleright}}\left(\mathsf{min}\{\,|N(y)\cap A|\,, |N(y)\cap B|\,\} - \frac{|N(y)\cap A\cap B|}{2}\right) \\
&\geqslant \left|N(x)\cap B_A^{\triangleright}\right| \geqslant |\{x\}, X|_X.
\end{aligned}$$

□

Lemma 4.6. *Let $(A, B) \in \vec{S}(X)$ such that $|A, B|_X > |(A, B)^{\triangleright}|_Y$. Then there either exists an $x \in (A\setminus B)\cup(B\setminus A)$ with $|\{x\}, X|_X \leqslant |A, B|_X$ such that $(A', B') := (A+x,\, B+x)$ has order $|A', B'|_X < |A, B|_X$, or there exists an $x \in A\cap B$ with $|\{x\}, X|_X \leqslant |A, B|_X$ such that either $(A', B') = (A-x, B)$ or $(A', B') = (A, B-x)$ has order $|A', B'|_X < |A, B|_X$.*

Proof. By the proof of Lemma 4.4, if $|A, B|_X > |A_B^{\triangleright}, B_A^{\triangleright}|_Y$, then

$$\begin{aligned}
&\sum_{x\in X}\mathsf{min}\{\,\left|N(x)\cap A_B^{\triangleright}\right|, \left|N(x)\cap B_A^{\triangleright}\right|\,\} \\
&\qquad < \sum_{a\in A}\left|N(a)\cap B_A^{\triangleright}\right| + \sum_{b\in B}\left|N(b)\cap A_B^{\triangleright}\right| - \sum_{x\in A\cap B}\frac{|N(x)|}{2}.
\end{aligned}$$

Thus, without loss of generality there needs to be an $x \in A$ such that

$$\left|N(x)\cap B_A^{\triangleright}\right| > \left|N(x)\cap A_B^{\triangleright}\right|.$$

Every such x is suitable for the x in the assumption by Lemma 4.5.

Now suppose that $x \in A\cap B$ and $\left|N(x)\cap B_A^{\triangleright}\right| > \left|N(x)\cap A_B^{\triangleright}\right|$. Then

$$\begin{aligned}
&|A\setminus\{x\}, B|_X \\
&= \sum_{y\in Y}\left(\mathsf{min}\{\,|N(y)\cap (A\setminus\{x\})|\,, |N(y)\cap B|\,\} - \frac{|N(y)\cap (A\cap B)\setminus\{x\}|}{2}\right)
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{y \in Y} \left(\mathsf{min}\{\, |N(y) \cap A| \,, |N(y) \cap B| \,\} - \frac{|N(y) \cap A \cap B|}{2} \right) \\
&\qquad - E(B_A^{\triangleright}, \{x\}) + \frac{|N(x)|}{2} \\
&< |A, B|_X.
\end{aligned}
$$

For every such $x \in X$ we have, since $x \in A \cap B$, that

$$
\begin{aligned}
&\mathsf{min}\{\, |N(y) \cap A| \,, |N(y) \cap B| \,\} - \frac{|N(y) \cap A \cap B|}{2} \\
&\geqslant |N(y) \cap \{x\}| - \frac{|N(y) \cap \{x\}|}{2} = \frac{|N(y) \cap \{x\}|}{2}.
\end{aligned}
$$

Hence,

$$
\begin{aligned}
|\{x\}, X|_X = \frac{|N(x)|}{2} &= \sum_{y \in Y} \frac{|N(y) \cap \{x\}|}{2} \\
&\leqslant \sum_{y \in Y} \left(\mathsf{min}\{\, |N(y) \cap A| \,, |N(y) \cap B| \,\} - \frac{|N(y) \cap A \cap B|}{2} \right) \\
&= |A, B|_X.
\end{aligned}
$$

□

We are now ready to prove Theorem 6:

Proof of Theorem 6. Both τ'' and $\tau \cap \vec{S}_k(X)$ are tangles of $\vec{S}_k(X)$, suppose that they are distinct. Let $(A, B) \in \tau$ be a separation of minimal order with the property that $(B, A) \in \tau''$ and let us assume further, that among all those separations (A, B) is chosen so that $A \cap B$ is as large as possible.

Suppose first that $|A, B| = \left|A_B^{\triangleright}, B_A^{\triangleright}\right|$ and let $(A', B') = ((A, B)^{\triangleright})^{\triangleleft}$. Then, by Lemma 4.4, we have that $A \subseteq A'$ and $B \subseteq B'$. Since $(B', A') \in \tau$, it follows that $(A', B') \neq (A, B)$ and so we can pick $x \in (A' \setminus A) \cup (B' \setminus B)$. Note that $|A' + x, B' + x| \leqslant |A, B|$ by Lemma 4.5. Thus, by the choice of (A, B), either $(A' + x, B' + x) \in \tau'' \cap \tau$ which implies that $x \in A' \setminus A$ or $(B' + x, A' + x) \in \tau'' \cap \tau$ which implies $x \in B' \setminus B$. In any case, since $|\{x\}, X|_X \leqslant |A, B|$ (again by Lemma 4.5) and $(\{x\}, X) \in \tau'' \cap \tau$ this contradicts the fact that τ'', respectively $\tau \cap \vec{S}_k(X)$ are tangles.

If, on the other hand, $|A, B| > \left|A_B^{\triangleright}, B_A^{\triangleright}\right|$ then, by Lemma 4.6 there either exists $x \in (A \setminus B) \cup (B \setminus A)$ such that for $(A', B') = (A + x,\, B + x)$ we have $|A', B'|_X < |A, B|_X$ and $|\{x\}, X|_X \leqslant |A, B|_X$, or there exists $x \in A \cap B$ such that for $(A', B') = (A - x, B)$ or $(A', B') = (A, B - x)$ we have $|A', B'|_X < |A, B|_X$ and $|\{x\}, X|_X \leqslant |A, B|_X$.

However, either of these cases again contradicts the fact that τ'', respectively $\tau \cap \vec{S}_k(X)$, is a tangle as $(\{x\}, X) \in \tau \cap \tau''$ and, by the choice of (A, B), either $(A', B') \in \tau \cap \tau''$ or $(B', A') \in \tau \cap \tau''$, but the respective orientations of $\{\, A, B \,\}$, $\{\, A', B' \,\}$, and $\{\, \{x\}, X \,\}$ together contradict (†). □

There is also a more exciting way to prove Theorem 5 and Theorem 6 indirectly, albeit at the cost of a slight increase in the factors on k. This is to view the tangles of the two partition classes as two different facets of tangles on the edge set of the bipartite graph. We give these proofs in the next section.

4.2. Tangles on the edges

We will show that the tangles on the sides of a bipartite graph are related to a special kind of tangles defined on the separations of the edges. So let us give the notation required for these intermediate tangles of the edges.

Denote the set of all set separations of E, the edge set of our bipartite graph, as $\vec{S}(E)$, and the set of the corresponding unoriented separations as $S(E)$. The following order function on the separations in $\vec{S}(E)$ is a natural variation on our previous order function for separations in $\vec{S}(X)$:

$$|C,D|_E := \sum_{v\in V}\left(\mathsf{min}\{\,|E(v)\cap C|\,,|E(v)\cap D|\,\} - \frac{|E(v)\cap C\cap D|}{2}\right),$$

where $E(v)$ denotes the set of incident edges of v. We will again assume that any $\vec{S}_k(E)$ we consider does not contain (E,E).

We say that an orientation τ of a subset $\vec{S}_k(E)$ of $\vec{S}(E)$ is a *tangle* of $\vec{S}_k(E)$, if τ is an orientation of $\vec{S}_k(E)$ with the following property:

$$\textit{There are no } (C_1,D_1),(C_2,D_2),(C_3,D_3)\in\tau \textit{ with } C_1\cup C_2\cup C_3=E. \qquad (\dagger_E)$$

Given a separation in $\vec{S}(X)$, it is pretty immediate how to obtain a separation in $\vec{S}(E)$ which is 'dual' to this separation: A separation (A,B) of X naturally defines a separation $(A,B)^E := (E(A),E(B))$ of E, where $E(A)$ denotes the set of all edges of G which have an end vertex in A. Note that $((A,B)^E)^* = (B,A)^E$.

The other way around is less obvious, but it will be necessary to associate to each separation in $\vec{S}(E)$ a separation in $\vec{S}(X)$ and $\vec{S}(Y)$. We will do so similarly to how we associated to each separation in $\vec{S}(Y)$ a separation in $\vec{S}(X)$. There we obtained, given a separation $(A,B)\in\vec{S}(Y)$, a separation in $\vec{S}(X)$ by asking for every vertex in X whether that vertex has more neighbours in A or in B. Similarly, we will now ask, given a separation (C,D) in $\vec{S}(E)$, for each vertex in X, whether more of the adjacent edges lie in C or in D. Formally, given a separation (C,D) of E, we obtain a separation $(C,D)^{\blacktriangleleft} := (C_D^{\blacktriangleleft}, D_C^{\blacktriangleleft})$ of X by defining

$$C_D^{\blacktriangleleft} = \{\, x\in X : |E(x)\cap C| \geqslant |E(x)\cap D|\,\}$$

and

$$D_C^{\blacktriangleleft} = \{\, x\in X : |E(x)\cap C| \leqslant |E(x)\cap D|\,\}.$$

This shifting operation preserves the partial order of separations in the following sense:

Lemma 4.7. *If* $(C,D) \leqslant (C',D')$*, then* $(C,D)^{\blacktriangleleft} \leqslant (C',D')^{\blacktriangleleft}$

Proof. If $(C,D) \leqslant (C',D')$, then $C \subseteq C'$ and $D \supseteq D'$. Thus, for $x \in X$, we have that $|E(x) \cap C| \leqslant |E(x) \cap C'|$ and $|E(x) \cap D| \geqslant |E(x) \cap D'|$.

Now if $x \in C_D^{\blacktriangleleft}$, then $|E(x) \cap C| \geqslant |E(x) \cap D|$ and thus

$$|E(x) \cap C'| \geqslant |E(x) \cap C| \geqslant |E(x) \cap D| \geqslant |E(x) \cap D'|\,,$$

hence $x \in C'^{\blacktriangleleft}_{D'}$. Similarly, if $x \in D'^{\blacktriangleleft}_{C'}$, then $|E(x) \cap D'| \geqslant |E(x) \cap C'|$ and thus

$$|E(x) \cap D| \geqslant |E(x) \cap D'| \geqslant |E(x) \cap C'| \geqslant |E(x) \cap C|\,,$$

hence $x \in C_D^{\blacktriangleleft}$. Thus $C_D^{\blacktriangleleft} \subseteq C'^{\blacktriangleleft}_{D'}$ and $D'^{\blacktriangleleft}_{C'} \subseteq D_C^{\blacktriangleleft}$, i.e., $(C,D)^{\blacktriangleleft} \leqslant (C',D')^{\blacktriangleleft}$. □

Symmetrically we can define a separation $(C,D)^{\blacktriangleright}$ of Y, but by the symmetry of the situation we will only ever need to talk about the map $(\cdot)^{\blacktriangleleft}$.

Unlike the shifting operations considered in the previous section, there is less of a symmetry here: The separation $(A,B)^E$ fully determines the separation (A,B), whereas the separation $(C,D)^{\blacktriangleleft}$ in some way 'compresses' the information in the separation (C,D) into a rough estimate. Generally there are multiple different separations (C,D) in $\vec{S}(E)$ for which the $(C,D)^{\blacktriangleleft}$ coincide, and so the operation $(\cdot)^{\blacktriangleleft}$ is not injective.

As with $(\cdot)^{\triangleright}$, this function induces a pull-back map: given a subset τ of $\vec{S}(X)$, we define

$$\tau_E := \{\, (C,D) : (C,D)^{\blacktriangleleft} \in \tau \,\}.$$

Note that, as $(E(A),E(B))^{\blacktriangleleft} = (A,B)$, the set of all the separations $(A,B)^E$ is a subset of τ_E.

For shifting in the other direction we take a slightly different notion. Given a tangle τ of $\vec{S}(E)$, let us define

$$\tau_X := \{\, (C,D)^{\blacktriangleleft} : (C,D) \in \tau \,\},$$

and let τ_Y be defined analogously. Note that this is a genuinely different way to move between tangles of $\vec{S}(E)$ and $\vec{S}(X)$; rather than 'pulling back' the tangle from $\vec{S}(X)$ to $\vec{S}(E)$ via the shift $(\cdot)^E$, giving rise to a set of separations

$${}^X\tau := \{\, (A,B) \in \vec{S}(X) : (A,B)^E \in \tau \,\},$$

we're 'pushing forward' via the shift $(\cdot)^{\blacktriangleleft}$.

We note that in this particular case, since, assuming the graph is connected, it is clear that $((A,B)^E)^{\blacktriangleleft} = (A,B)$, we have that $\tau_X \supseteq {}^X\tau$ and so, since ${}^X\tau$ is automatically a partial orientation of $\vec{S}(X)$, if the restriction of τ_X to some lower order is a tangle, then the restriction of ${}^X\tau$ to the same order will also satisfy (†). In particular, working with this definition results in slightly stronger results than working with ${}^{X}\tau$, however the main purpose of this change is that it will result in slightly simpler proofs, see for example Corollary 4.14.

We will show that, for a given tangle τ of $\vec{S}_{4k}(X)$, the set

$$\tau_{E,k} := \tau_{E,k}$$

actually is a tangle of $\vec{S}_k(E)$ and dually, that if τ is a tangle of $\vec{S}_{2k}(E)$, then

$$\tau_{X,k} := \tau_{X,k}$$

is a tangle of $\vec{S}_k(X)$. We will then be able to use this to obtain proofs of Theorem 5 and Theorem 6 from the symmetry between X and Y.

Proposition 4.8. *The order function $|\cdot|_E$ is submodular.*

Proof. As in the proof of Proposition 4.1, it is enough to show that, for every $v \in V$, the function

$$|C, D|_v := \mathsf{min}\{\, |E(v) \cap C|\,, |E(v) \cap D|\,\} - \frac{|E(v) \cap C \cap D|}{2}$$

is submodular, as clearly $|C, D|_E = \sum_{v \in V} |C, D|_v$. Now fix some v in V. For $F \subseteq E$ we denote $N_F := |E(v) \cap F|$.

Take separations (C_1, D_1) and (C_2, D_2) in $\vec{S}(E)$ and suppose without loss of generality that $N_{C_i} \leqslant N_{D_i}$. Let $C_i' := C_i \setminus D_i$, $D_i' := D_i \setminus C_i$ and $F_i := C_i \cap D_i$. Then

$$|C_1 \cap C_2, D_1 \cup D_2|_v = N_{C_1' \cap C_2'} + \frac{1}{2}(N_{F_1 \cap F_2} + N_{F_1 \cap C_2'} + N_{C_1' \cap F_2}),$$

and

$$|C_1 \cup C_2, D_1 \cap D_2|_v = \mathsf{min}\{\, N_{C_1' \cup C_2'}, N_{D_1' \cap D_2'}\,\} + \frac{1}{2}(N_{F_1 \cap F_2} + N_{F_1 \cap D_2'} + N_{D_1' \cap F_2}).$$

Summing these two, we get

$$\begin{aligned}
&|C_1 \cap C_2, D_1 \cup D_2|_v + |C_1 \cup C_2, D_1 \cap D_2|_v \\
&= N_{C_1' \cap C_2'} + \frac{1}{2}(N_{F_1 \cap F_2} + N_{F_1 \cap C_2'} + N_{C_1' \cap F_2}) \\
&\quad + \mathsf{min}\{\, N_{C_1' \cup C_2'}, N_{D_1' \cap D_2'}\,\} + \frac{1}{2}(N_{F_1 \cap F_2} + N_{F_1 \cap D_2'} + N_{D_1' \cap F_2}) \\
&\leqslant N_{C_1' \cap C_2'} + N_{C_1' \cup C_2'} \\
&\quad + \frac{1}{2}(N_{F_1 \cap F_2} + N_{F_1 \cap C_2'} + N_{F_1 \cap D_2'} + N_{F_1 \cap F_2} + N_{C_1' \cap F_2} + N_{D_1' \cap F_2}) \\
&= N_{C_1'} + N_{C_2'} + \frac{1}{2}(N_{F_1} + N_{F_2}) = |C_1, D_1|_v + |C_2, D_2|_v\,.
\end{aligned}$$

Similarly

$$|C_1 \cap D_2, D_1 \cup C_2|_v = \mathsf{min}\{\, N_{C_1' \cap D_2'}, N_{D_1' \cup C_2'} \,\} + \frac{1}{2}(N_{F_1 \cap F_2} + N_{F_1 \cap D_2'} + N_{C_1' \cap F_2}),$$

and

$$|C_1 \cup D_2, D_1 \cap C_2|_v = \mathsf{min}\{\, N_{C_1' \cup D_2'}, N_{D_1' \cap C_2'} \,\} + \frac{1}{2}(N_{F_1 \cap F_2} + N_{F_1 \cap C_2'} + N_{D_1' \cap F_2}).$$

Summing these two, we get

$$\begin{aligned}
&|C_1 \cap D_2, D_1 \cup C_2|_v + |C_1 \cup D_2, D_1 \cap C_2|_v \\
&= \mathsf{min}\{\, N_{C_1' \cap D_2'}, N_{D_1' \cup C_2'} \,\} + \frac{1}{2}(N_{F_1 \cap F_2} + N_{F_1 \cap D_2'} + N_{C_1' \cap F_2}) \\
&\quad + \mathsf{min}\{\, N_{C_1' \cup D_2'}, N_{D_1' \cap C_2'} \,\} + \frac{1}{2}(N_{F_1 \cap F_2} + N_{F_1 \cap C_2'} + N_{D_1' \cap F_2}) \\
&\leqslant N_{C_1' \cap D_2'} + N_{C_2' \cap D_1'} \\
&\quad + \frac{1}{2}(N_{F_1 \cap F_2} + N_{F_1 \cap C_2'} + N_{F_1 \cap D_2'} + N_{F_1 \cap F_2} + N_{C_1' \cap F_2} + N_{D_1' \cap F_2}) \\
&\leqslant N_{C_1'} + N_{C_2'} + \frac{1}{2}(N_{F_1} + N_{F_2})\, |C_1, D_1|_v + |C_2, D_2|_v\,.
\end{aligned}$$

Thus $|\cdot|_v$ is a submodular function, and so is $|\cdot|_E$. □

However, unlike for the correspondence between $|\cdot|_X$ and $|\cdot|_Y$, we will no longer be able to show that the order of the shift of a separation is non-increasing, instead we will only be able to show that, when shifting from a separation of the vertices to the corresponding separation of the edges, we can bound how much the order increases. More precisely, simple calculations show that:

Proposition 4.9. *Given a separation (A, B) of X, we have $|A, B|_X \leqslant |(A, B)^E|_E$ and $|(A, B)^E|_E \leqslant 2|A, B|_X$.*

Proof. For the first statement we note that

$$\begin{aligned}
|A, B|_X &= \sum_{y \in Y} \left(\mathsf{min}\{\, |N(y) \cap A| \,, |N(y) \cap B| \,\} - \frac{|N(y) \cap A \cap B|}{2} \right) \\
&= \sum_{y \in Y} \left(\mathsf{min}\{\, |E(y) \cap E(A)| \,, |E(y) \cap E(B)| \,\} - \frac{|E(y) \cap E(A) \cap E(B)|}{2} \right) \\
&\leqslant \sum_{v \in V} \left(\mathsf{min}\{\, |E(v) \cap E(A)| \,, |E(v) \cap E(B)| \,\} - \frac{|E(v) \cap E(A) \cap E(B)|}{2} \right) \\
&= |E(A), E(B)|_E.
\end{aligned}$$

For the second statement we observe that, for $x \in X$ we have that

$$\mathsf{min}\{\, |E(x) \cap E(A)| \,, |E(x) \cap E(B)| \,\} = |E(x) \cap E(A) \cap E(B)|$$

and thus

$$\sum_{x\in X}\left(\mathsf{min}\{\,|E(x)\cap E(A)|\,,|E(x)\cap E(B)|\,\}-\frac{|E(x)\cap E(A)\cap E(B)|}{2}\right)$$
$$=\frac{|E(A)\cap E(B)|}{2}.$$

As clearly $|A,B|_X \geqslant \frac{|E(A)\cap E(B)|}{2}$, it follows that

$$\begin{aligned}
&|E(A),E(B)|_E\\
&=\sum_{v\in V}\left(\mathsf{min}\{\,|E(v)\cap E(A)|\,,|E(v)\cap E(B)|\,\}-\frac{|E(v)\cap E(A)\cap E(B)|}{2}\right)\\
&=\sum_{x\in X}\left(\mathsf{min}\{\,|E(x)\cap E(A)|\,,|E(x)\cap E(B)|\,\}-\frac{|E(x)\cap E(A)\cap E(B)|}{2}\right)\\
&\quad+\sum_{y\in Y}\left(\mathsf{min}\{\,|E(y)\cap E(A)|\,,|E(y)\cap E(B)|\,\}-\frac{|E(y)\cap E(A)\cap E(B)|}{2}\right)\\
&=\frac{|E(A)\cap E(B)|}{2}+\sum_{y\in Y}\left(\mathsf{min}\{\,|N(y)\cap A|\,,|N(y)\cap B|\,\}-\frac{|N(y)\cap A\cap B|}{2}\right)\\
&\leqslant 2|A,B|_X.
\end{aligned}$$

□

For $(\cdot)^{◀}$ on the other hand, we will be able to show that this is a non-increasing operation:

Lemma 4.10. *Let (C,D) be a separation of E, then $|C,D|_E \geqslant |C^{◀}_D, D^{◀}_C|_X$.*

For the proof of Lemma 4.10 we will need to carefully analyse how we can 'locally' change a separation in $\vec{S}(E)$ without changing the shift. Recall that, given a separation (A,B) in $\vec{S}(X)$, there are other separations apart from $(A,B)^E$ in $\vec{S}(E)$ which still shift to (A,B). So, in order to prove Lemma 4.10 we will analyse what these different separations of E inducing the same separation (A,B) of X look like. For this, we will show which 'local', i.e., single-edge, changes we can make to a given separation (C,D) to bring it closer to one of the type $(A,B)^E$, without increasing its order.

So, let us start analysing these 'local' changes. Firstly, in the next lemma we show that we can move a single edge from C to D without increasing the order of (C,D) or changing its shift $(C,D)^{◀}$, if at the end vertex in X of that edge there are fewer incident edges in C than in D.

Lemma 4.11. *Let (C,D) be a separation of E and let $e\in E$ be incident with $C^{◀}_D \setminus D^{◀}_C$. Then $|C+e,D-e|_E \leqslant |C,D|_E$ and $(C+e,D-e)^{◀} = (C,D)^{◀}$.*

Proof. Let $e=vw$. We observe that, since $v\in C^{◀}_D\setminus D^{◀}_C$, we have

$$\mathsf{min}\{\,|E(v)\cap C|\,,|E(v)\cap D|\,\}-\frac{|E(v)\cap C\cap D|}{2}$$
$$=|E(v)\cap D|-\frac{|E(v)\cap C\cap D|}{2}$$

$$\geqslant |E(v) \cap (D-e)| - \frac{|E(v) \cap (C+e) \cap (D-e)|}{2} + 1,$$

and

$$\begin{aligned}
&\mathsf{min}\{\, |E(w) \cap C| \,, |E(w) \cap D| \,\} - \frac{|E(w) \cap C \cap D|}{2} \\
&\geqslant \mathsf{min}\{\, |E(w) \cap (D-e)| \,, |E(w) \cap (C+e)| \,\} - \frac{|E(w) \cap (C+e) \cap (D-e)|}{2} - 1.
\end{aligned}$$

Thus, $|C-e, D+e|_E \leqslant |C,D|_E$. Moreover, v lies in $C^{\blacktriangleleft}_D \setminus D^{\blacktriangleleft}_C$ and therefore also in $(C+e)^{\blacktriangleleft}_{D-e} \setminus (D-e)^{\blacktriangleleft}_{C+e}$, and thus $(C,D)^{\blacktriangleleft} = (C+e, D-e)^{\blacktriangleleft}$. $\square$

Using this, we can show that from within the set of those separations of E with the same shift (A,B), we can always find some (C,D) of minimal order which is 'close' to $(A,B)^E$, in the sense that every edge incident with $A \setminus B$ is contained in $C \setminus D$ and every edge incident with $B \setminus A$ is contained in $D \setminus C$:

Lemma 4.12. *Let $(C,D) \in \vec{S}(E)$. Then there exists a separation (C',D') of E with $|C',D'|_E \leqslant |C,D|_E$ and $(C',D')^{\blacktriangleleft} = (C,D)^{\blacktriangleleft}$ such that every edge e incident with $C^{\blacktriangleleft}_D \setminus D^{\blacktriangleleft}_C$ lies in $C' \setminus D'$ and every edge incident with $D^{\blacktriangleleft}_C \setminus C^{\blacktriangleleft}_D$ lies in $D' \setminus C'$.*

Proof. Suppose (C',D') is chosen so that $|C',D'|_E \leqslant |C,D|_E$ and $(C',D')^{\blacktriangleleft} = (C,D)^{\blacktriangleleft}$, and so that there are as few edges as possible incident with $C^{\blacktriangleleft}_D \setminus D^{\blacktriangleleft}_C$ which do not lie in $C' \setminus D'$ and as few edges as possible incident with $D^{\blacktriangleleft}_C \setminus C^{\blacktriangleleft}_D$ which do not lie in $D' \setminus C'$.

Suppose that there exists some such edge e incident with $C^{\blacktriangleleft}_D \setminus D^{\blacktriangleleft}_C$ which does not lie in $C' \setminus D'$ or some such edge e incident with $D^{\blacktriangleleft}_C \setminus C^{\blacktriangleleft}_D$ which does not lie in $D' \setminus C'$. Let us assume we are in the former case, as the argument in the latter case is identical.

Since $(C',D')^{\blacktriangleleft} = (C,D)^{\blacktriangleleft}$, by Lemma 4.11 we could then consider the separation $(C'+e, D'-e)$ which must then satisfy $(C'+e, D'-e)^{\blacktriangleleft} = (C',D')^{\blacktriangleleft}$ and $|C'+e, D'-e| \leqslant |C',D'| \leqslant |C,D|$, contradicting the choice of (C',D'). $\square$

This observation enables us to perform the necessary calculations to prove Lemma 4.10.

Proof of Lemma 4.10. By Lemma 4.12 we may suppose that every edge incident with $C^{\blacktriangleleft}_D \setminus D^{\blacktriangleleft}_C$ lies in $C \setminus D$ and every edge incident with $D^{\blacktriangleleft}_C \setminus C^{\blacktriangleleft}_D$ lies in $D \setminus C$.

In this case, we can calculate $|C,D|_E$ as follows.

$$\begin{aligned}
&|C,D|_E \\
&= \sum_{v \in V} \left(\mathsf{min}\{\, |E(v) \cap C| \,, |E(v) \cap D| \,\} - \frac{|E(v) \cap C \cap D|}{2} \right) \\
&= \sum_{v \in Y} \left(\mathsf{min}\{\, |E(v) \cap C| \,, |E(v) \cap D| \,\} - \frac{|E(v) \cap C \cap D|}{2} \right) + \sum_{v \in C^{\blacktriangleleft}_D \cap D^{\blacktriangleleft}_C} \frac{1}{2} \, |E(v)|
\end{aligned}$$

$$
\begin{aligned}
&\geqslant \sum_{v\in Y}\left(\mathsf{min}\{\left|N(v)\cap C_D^{\blacktriangleleft}\right|,\left|N(v)\cap D_C^{\blacktriangleleft}\right|\}-\frac{1}{2}\left|N(v)\cap C_D^{\blacktriangleleft}\cap D_C^{\blacktriangleleft}\right|\right)\\
&\qquad -\frac{1}{2}\left|E(C_D^{\blacktriangleleft})\cap E(D_C^{\blacktriangleleft})\right|+\sum_{v\in C_D^{\blacktriangleleft}\cap D_C^{\blacktriangleleft}}\frac{1}{2}\left|N(v)\right|\\
&\geqslant \sum_{v\in Y}\left(\mathsf{min}\{\left|N(v)\cap C_D^{\blacktriangleleft}\right|,\left|N(v)\cap D_C^{\blacktriangleleft}\right|\}-\frac{\left|N(v)\cap C_D^{\blacktriangleleft}\cap D_C^{\blacktriangleleft}\right|}{2}\right)\\
&=|C_D^{\blacktriangleleft},D_C^{\blacktriangleleft}|_X
\end{aligned}
$$
□

Analysing local changes will also play a crucial role in showing that, given a tangle τ in $\vec{S}(E)$, the restriction of τ_X to a lower order is actually an orientation. For this we will need to make sure that separations obtained from one another by local changes cannot be oriented differently in τ_E. However, whereas Lemma 4.11 allows us to move certain edges from $C \setminus D$ to $D \setminus C$ without changing the shift or increasing the order, for showing that the restriction of τ_X to a lower order actually is an orientation we will need to analyse a different type of local change.

More precisely, the next lemma will allow us to move certain edges from $D \setminus C$ (or, symmetrically, $C \setminus D$), to $C \cap D$ without increasing the order. Such an operation might change the shift of a separation, but, by Lemma 4.7, it does so only in a controlled way: moving an edge from $D \setminus C$ to $D \cap C$ will only result in a shift that is larger, in the sense of the partial order on the separation system, than the shift of the original (C, D). Moreover, such a local change does not change the way a separation is oriented by a tangle.

Lemma 4.13. *Let (C, D) be a separation of E and let $e \in E$ be incident with $C_D^{\blacktriangleleft}$. Then $|C + e, D|_E \leqslant |C, D|_E$ and $(C + e, D)^{\blacktriangleleft} \geqslant (C, D)^{\blacktriangleleft}$.*

Proof. If $e \in C$, then there is nothing to show, so suppose $e \in D \setminus C$ and let $e = vw$. We observe that, since $v \in C_D^{\blacktriangleleft}$, we have

$$
\begin{aligned}
&\mathsf{min}\{\,|E(v)\cap C|\,,|E(v)\cap D|\,\}-\frac{|E(v)\cap C\cap D|}{2}\\
&=|E(v)\cap D|-\frac{|E(v)\cap C\cap D|}{2}\\
&=|E(v)\cap D|-\frac{|E(v)\cap (C+e)\cap D|}{2}+\frac{1}{2}
\end{aligned}
$$

and

$$
\begin{aligned}
&\mathsf{min}\{\,|E(w)\cap C|\,,|E(w)\cap D|\,\}-\frac{|E(w)\cap C\cap D|}{2}\\
&\qquad\geqslant \mathsf{min}\{\,|E(w)\cap (C+e)|\,,|E(w)\cap D|\,\}-\frac{|E(w)\cap (C+e)\cap D|}{2}-\frac{1}{2}.
\end{aligned}
$$

Thus, $|C + e, D|_E \leqslant |C, D|_E$. We have $(C, D)^{\blacktriangleleft} \leqslant (C + e, D)^{\blacktriangleleft}$ by Lemma 4.7. □

We now have all the ingredients at hand needed to show that the shift of a tangle, restricted to an appropriate order, is still a tangle. Let us start by considering the shift τ_X of a tangle τ in $\vec{S}(E)$.

Theorem 7. *If τ is a tangle of $\vec{S}_{2k}(E)$, then $\tau_{X,k}$ is a tangle of $\vec{S}_k(X)$.*

Proof. We first note that the set $\tau_{X,k}$ contains at least one of (A,B) and (B,A) for every separation $(A,B) \in \vec{S}_k(X)$. Indeed, by Proposition 4.9 $|(A,B)^E|_E \leqslant 2|A,B|_X$, and so since τ is a tangle of $\vec{S}_{2k}(E)$ either $(A,B)^E \in \tau$ or $(B,A)^E \in \tau$.

Let us now show that for no separation $\{A,B\}$ do we have both (A,B) and (B,A) in $\tau_{X,k}$. Suppose otherwise, then τ contains separations (C_1,D_1) and (C_2,D_2) such that $(C_1,D_1)^{\blacktriangleleft} = (A,B)$ and $(C_2,D_2)^{\blacktriangleleft} = (B,A)$.

Note that, by Proposition 4.9, we have that $|E(A),E(B)|_E \leqslant 2|A,B|_X < 2k$, hence $(A,B)^E \in \tau$ or $(B,A)^E \in \tau$. As $(E(A),E(B))^{\blacktriangleleft} = (A,B)$, we may suppose without loss of generality that either $(C_1,D_1) = (A,B)^E$ or $(C_2,D_2) = (B,A)^E$. We suppose the former one, the latter case is similar.

Now pick a separation $(C,D) \in \tau$ so that $(C,D)^{\blacktriangleleft} \geqslant (C_2,D_2)^{\blacktriangleleft} = (B,A)$ and the set $(D_1 \cap D) \setminus (C_1 \cup C)$ is as small as possible. Then, since τ satisfies $(\dagger_E)$, we have $C_1 \cup C \neq E$. Hence there exists some edge $e \in (D_1 \cap D) \setminus (C_1 \cup C)$.

Let x be the end vertex of e in X. Note that $e \in E(B) \setminus E(A)$ since $E(A) = C_1$ and $e \notin C_1$. Thus, $x \in B \setminus A$. Moreover, as $B = C_2{}^{\blacktriangleleft}_{D_2} \subseteq C^{\blacktriangleleft}_D$, we have that $x \in C^{\blacktriangleleft}_D$. Thus, e is incident with $C^{\blacktriangleleft}_D$.

Consequently, we can apply Lemma 4.13 to get that $|C+e,D| \leqslant |C,D|$. Thus, τ orients $(C+e,D)$ and therefore $(C+e,D) \in \tau$, as $(D,C+e) \in \tau$ would contradict $(\dagger_E)$ because of $(C,D) \in \tau$ and $D \cup C = E$.

But this implies that $(C+e,D) \in \tau$ is a better choice for (C,D), since $(C+e,D)^{\blacktriangleleft} \geqslant (C,D)^{\blacktriangleleft}$ by Lemma 4.7, and

$$(D_1 \cap D) \setminus (C_1 \cup C) \supsetneq (D_1 \cap D) \setminus (C_1 \cup (C+e)),$$

as $e \in (D_1 \cap D) \setminus (C_1 \cup C)$.

Thus, $\tau_{X,k}$ is an orientation of $\vec{S}_k(X)$. That $\tau_{X,k}$ satisfies the tangle property $(\dagger)$ now follows like this: if (A_1,B_1), (A_2,B_2), (A_3,B_3) were a triple of separations in $\tau_{X,k}$ contradicting the tangle property $(\dagger)$, then τ would need to orient $(A_1,B_1)^E$, $(A_2,B_2)^E$ and $(A_3,B_3)^E$ by Proposition 4.9. By the above observation, τ then orients them as $(A_1,B_1)^E$, $(A_2,B_2)^E$ and $(A_3,B_3)^E$, since $\tau_{X,k}$ does not contain any $(((A_i,B_i)^E)^*)^{\blacktriangleleft} = (B_i,A_i)$. However, the three separations $(A_1,B_1)^E$, $(A_2,B_2)^E$ and $(A_3,B_3)^E$ in τ then contradict the tangle property $(\dagger_E)$, as every edge in E is incident with at least one of the sets A_1, A_2, and A_3. $\square$

A similar conclusion holds for the shift τ_E of a tangle τ of $\vec{S}(X)$.

Theorem 8. *Given a tangle τ of $\vec{S}_{4k}(X)$, then $\tau_{E,k}$ is a tangle of $\vec{S}_k(E)$.*

Proof. By Lemma 4.10, given some separation $(C,D) \in \vec{S}_k(E)$, we have that $|(C,D)^{◀}|_X \leqslant |C,D|_E$, thus τ contains exactly one of $(C,D)^{◀}$ and $((C,D)^{◀})^* = (D,C)^{◀}$, and consequently $\tau_{E,k}$ contains exactly one of (C,D) and (D,C), i.e., $\tau_{E,k}$ is an orientation of $\vec{S}_k(E)$.

So, it remains to show that $\tau_{E,k}$ satisfies the tangle property $(\dagger_E)$. Let us suppose for a contradiction that there is some set

$$\{ (C_1,D_1),(C_2,D_2),(C_3,D_3) \} \subseteq \tau_{E,k}$$

such that $C_1 \cup C_2 \cup C_3 = E$.

Let $(A_i,B_i) = (C_i,D_i)^{◀}$ for each $i = 1,2,3$. Then, since $(A_i,B_i) \in \tau$ for each i, and τ is a tangle, it follows that the set $Z = X \setminus (A_1 \cup A_2 \cup A_3)$ is non-empty.

Since $Z \subseteq B_i = D_{i\,C_i}^{◀}$ for each i, we have that $|E(z) \cap D_i| \geqslant |E(z) \cap C_i|$ for all $z \in Z$ and $i = 1,2,3$. However, since $C_1 \cup C_2 \cup C_3 = E$,

$$\begin{aligned}
\sum_{i=1}^{3} |C_i,D_i|_E &= \sum_{i=1}^{3}\sum_{v\in V}\left(\mathsf{min}\{\,|E(v)\cap C_i|\,,|E(v)\cap D_i|\,\} - \frac{1}{2}\,|E(v)\cap C_i \cap D_i|\right)\\
&\geqslant \sum_{i=1}^{3}\sum_{z\in Z}\left(\mathsf{min}\{\,|E(z)\cap C_i|\,,|E(z)\cap D_i|\,\} - \frac{1}{2}\,|E(z)\cap C_i \cap D_i|\right)\\
&= \sum_{z\in Z}\sum_{i=1}^{3}\left(|E(z)\cap C_i| - \frac{1}{2}\,|E(z)\cap C_i \cap D_i|\right)\\
&\geqslant \sum_{z\in Z} d(z)/2 = \frac{|E(Z,Y)|}{2}.
\end{aligned}$$

As $|C_i,D_i|_E < k$ for every $i = 1,2,3$, this gives us $|E(Z,X)| < 6k$ and thus

$$|Z,X|_X = \frac{|E(Z,Y)|}{2} < 3k.$$

Hence, τ needs to orient (Z,X). As $(X,Z) \in \tau$ would contradict $(\dagger)$, it follows that $(Z,X) \in \tau$.

Finally, since $|A_3,B_3|_X \leqslant |C_3,D_3|_E < k$ by Lemma 4.10, we can conclude by submodularity, that

$$|A_3 \cup Z, B_3 \cap X|_X \leqslant |A_3,B_3|_X + |Z,X| < 4k.$$

Hence, it follows that τ needs to orient $(A_3 \cup Z, B_3)$ and as $(A_3,B_3) \in \tau$ it follows from $(\dagger)$ that $(A_3 \cup Z, B_3) \in \tau$, as $A_3 \cup B_3 = X$. However, then

$$\{ (A_1,B_1),(A_2,B_2),(A_3 \cup Z,B_3) \} \subseteq \tau$$

and $A_1 \cup A_2 \cup (A_3 \cup Z) = X$, contradicting $(\dagger)$. □

Corollary 4.14. *Let τ be a tangle of $\vec{S}_{8k}(E)$, then*

$$\tau'' := \left(\tau_{X,4k}\right)_{E,k}$$

is a subset of τ. Similarly, let τ' be a tangle of $\vec{S}_{8k}(X)$, then

$$\tau''' := \left(\tau'_{E,2k}\right)_{X,k}$$

is a subset of τ'.

Proof. By Theorem 7, τ'' is a tangle of $\vec{S}_k(E)$. Now, given any separation $(C,D) \in \vec{S}_k(E) \cap \tau$, we have that $(C,D)^{\blacktriangleleft} \in \tau_{X,4k}$ and thus (C,D) is in τ''. As τ'' is an orientation of $\vec{S}_k(E)$, we then have that $\tau'' \subseteq \tau$.

For the second part we note that, by Theorem 8, τ''' is a tangle of $\vec{S}_k(X)$. Given $(A,B) \in \vec{S}_k(X) \cap \tau'$ we have, since $((A,B)^E)^{\blacktriangleleft} = (A,B)$, that $(A,B)^E$ is in $\tau'_{E,2k}$ and thus (A,B) is in τ'''. As τ''' is an orientation of $\vec{S}_k(X)$, we then have that $\tau''' \subseteq \tau'$. $\square$

Putting these together, we obtain versions of Theorem 5 and Theorem 6, with slightly worse factors:

Corollary 4.15. *Let τ be a tangle of $\vec{S}_{8k}(X)$. Then $\tau' := {}^{\triangleright}_{k}\tau$ is a tangle of $\vec{S}_k(Y)$.*

Proof. It is easy to see that $\tau' = (\tau_{E,2k})_{Y,k}$ which is a tangle by Theorem 8 and Theorem 7. $\square$

Corollary 4.16. *Let τ be a tangle of $\vec{S}_{64k}(X)$, let $\tau' = {}^{\triangleright}_{8k}\tau$, and let $\tau'' = {}^{\triangleleft}_{k}\tau'$, then $\tau'' \subseteq \tau$.*

Proof. Consider $\tau_{E,16k}$. By Corollary 4.15, we have that

$$\tau' = (\tau_{E,16k})_{Y,8k}.$$

Moreover, again by Corollary 4.15 we have that

$$\tau'' = (\tau'_{E,2k})_{X,k}.$$

But now, by Corollary 4.14, we note that

$$\left(\left(\tau_{E,16k}\right)_{Y,8k}\right)_{E,2k} \subseteq \tau_{E,16k}$$

and thus,

$$\tau'' = \left(\left(\left(\tau_{E,16k}\right)_{Y,8k}\right)_{E,2k}\right)_{X,k} \subseteq \left(\tau_{E,16k}\right)_{X,k}.$$

Again by Corollary 4.14 we have that

$$\left(\tau_{E,16k}\right)_{X,k} \subseteq \tau,$$

which shows the claim. $\square$

4.3. Variations and generalizations

A natural question to consider at this point is how much these results depend on the very specific set up we have here.

For example, whilst we considered a very specific type of tangle, there are other types of ‘tangle-like’ clusters which one might wish to consider. Perhaps the most general condition one could consider here would be that of a *regular profile*. Recall that, for systems of set separations, a *profile* is a consistent orientation of the separation system which satisfies property (P), i.e., it does not contain any triple of separations of the form

$$\{ (A_1, B_1), (A_2, B_2), (B_1 \cap B_2, A_1 \cup A_2) \}.$$

A *regular* profile is one which does not contain any cosmall separations, that is, a separation (V, B) where V is the underlying set and $B \subseteq V$. Tangles are regular profiles, but regular profiles model a broader class of clusters.

Similar statements as in Theorem 5 and 6 can be shown via similar arguments for regular profiles, with only slight adjustments to the proofs and reusing our lemmas:

Theorem 9. *Let P be a regular profile of $\vec{S}_{3k}(X)$, then $P' := {}^{\triangleright}_{k}P$ is a regular profile of $\vec{S}_k(Y)$.*

Proof. We first note that P' is an orientation of $\vec{S}_k(Y)$. To show that P' is consistent, suppose for the contrary that there are some separations $(C_1, D_1), (C_2, D_2) \in P'$ such that $(D_1, C_1) \leqslant (C_2, D_2)$ and let $(C_i, D_i)^{\triangleleft} = (A_i, B_i)$. Then, since $D_1 \subseteq C_2$ and $D_2 \subseteq C_1$ we have that $(B_1, A_1) \leqslant (A_2, B_2)$ and thus $(A_1, B_1), (A_2, B_2) \in P$ contradicts the consistency of P.

So, it remains to show that P' is a regular profile. Let us suppose for a contradiction that there is some set $\{ (C_1, D_1), (C_2, D_2), (C_3, D_3) \} \subseteq P'$ such that $(C_1, D_1) \vee (C_2, D_2) = (D_3, C_3)$. Let $(A_i, B_i) = (C_i, D_i)^{\triangleleft}$ for each $i = 1, 2, 3$ and let $Z = B_1 \cap B_2 \cap B_3$.

Since $Z \subseteq B_i$ for each i, we have that $|N(x) \cap D_i| \geqslant |N(x) \cap C_i|$ for all $x \in Z$ and $i = 1, 2, 3$. However, since $C_1 \cup C_2 \cup C_3 = Y$,

$$\begin{aligned}
\sum_{i=1}^{3} |C_i, D_i|_Y &= \sum_{i=1}^{3} \sum_{x \in X} \Big(\mathsf{min}\{ |N(x) \cap C_i| , |N(x) \cap D_i| \} - \frac{|N(x) \cap C_i \cap D_i|}{2} \Big) \\
&\geqslant \sum_{i=1}^{3} \sum_{x \in Z} \Big(\mathsf{min}\{ |N(x) \cap C_i| , |N(x) \cap D_i| \} - \frac{|N(x) \cap C_i \cap D_i|}{2} \Big) \\
&\geqslant \sum_{x \in Z} \sum_{i=1}^{3} \Big(|N(x) \cap C_i| - \frac{|N(x) \cap C_i \cap D_i|}{2} \Big) \\
&\geqslant \sum_{z \in Z} d(z)/2 = \frac{|E(Z, Y)|}{2}.
\end{aligned}$$

Since $|C_i, D_i|_Y < k$ for every $i = 1, 2, 3$, we have $|E(Z, Y)| < 6k$ and thus

$$|Z, X|_X = E(Z, Y)/2 \leqslant 3k.$$

Hence $(Z, X) \in P$ since P is a regular profile of $\vec{S}_{3k}(X)$.

Moreover, P contains $(A_1, B_1) \vee (A_2, B_2)$ as by submodularity

$$|(A_1, B_1) \vee (A_2, B_2)|_X \leqslant |A_1, B_1|_X + |A_2, B_2|_X < 2k.$$

Then also, $(A_1, B_1) \vee (A_2, B_2) \vee (A_3, B_3) \in P$ as

$$|(A_1, B_1) \vee (A_2, B_2) \vee (A_3, B_3)|_X \leqslant |(A_1, B_1) \vee (A_2, B_2)|_X + |A_3, B_3|_X < 3k.$$

However, $(A_1, B_1) \vee (A_2, B_2) \vee (A_3, B_3) = (A_1 \cup A_2 \cup A_3, Z)$ and, since $Z = B_1 \cap B_2 \cap B_3$, we have that $A_1 \cup A_2 \cup A_3 \cup Z = X$ and thus $(A_1 \cup A_2 \cup A_3, Z) \vee (Z, X) = (X, Z)$, which contradicts the fact that P is a profile.

The profile P' is regular, since if $(Y, C) \in P'$, then $(Y, C)^{\triangleright} = (X, C_Y^{\triangleright})$ is a cosmall separation in P, which contradicts the regularity of P. □

Theorem 10. *Let P be a regular profile of $\vec{S}_{9k}(X)$, let $P' = {}^{\triangleright}_{k}P$, and let $P'' = {}^{\triangleleft}_{k}P'$, then $P'' \subseteq P$.*

Proof. Both P'' and $P \cap \vec{S}_k(X)$ are regular profiles of $\vec{S}_k(X)$. Let us suppose that they are distinct. Let $(A, B) \in P$ be a separation of minimal order with the property that $(B, A) \in P''$ and let us assume further, that among all those separations (A, B) is chosen so that $A \cap B$ is as large as possible. Suppose first that $|A, B| = \left|A_B^{\triangleright}, B_A^{\triangleright}\right|$.

Let $(A', B') = ((A, B)^{\triangleright})^{\triangleleft}$, by Lemma 4.4, we have that $A \subseteq A'$ and $B \subseteq B'$. Since $(B', A') \in P$, we have $(A', B') \neq (A, B)$, and so we can pick $x \in (A' \setminus A) \cup (B' \setminus B)$. Note that $|A' + x, B' + x| \leqslant |A, B|$ by Lemma 4.5. Thus, by the choice of (A, B), either $(A' + x, B' + x) \in P'' \cap P$ which implies that $x \in A' \setminus A$ or $(B' + x, A' + x) \in P'' \cap P$ which implies $x \in B' \setminus B$. In any case, since $|\{x\}, X|_X \leqslant |A, B|$ (again by Lemma 4.5) and $(\{x\}, X) \in P'' \cap P$ this contradicts the fact that P'' and, respectively, $P \cap \vec{S}_k(X)$ are profiles.

If on the other hand $|A, B| > \left|A_B^{\triangleright}, B_A^{\triangleright}\right|$ then, by Lemma 4.6 there either exists $x \in (A \setminus B) \cup (B \setminus A)$ such that for $(A', B') = (A + x,\, B + x)$ we have $|A', B'|_X < |A, B|_X$ and $|\{x\}, X|_X \leqslant |A, B|_X$, or there exists $x \in A \cap B$ such that for $(A', B') = (A - x, B)$ or $(A', B') = (A, B - x)$ we have $|A', B'|_X < |A, B|_X$ and $|\{x\}, X|_X \leqslant |A, B|_X$.

However, either of these cases again contradicts the fact that P'' or, respectively, $P \cap \vec{S}_k(X)$, is a profile as $(\{x\}, X) \in P \cap P''$ and, by the choice of (A, B), either $(A', B') \in P \cap P''$ or $(B', A') \in P \cap P''$, and $\{A, B\}$, $\{A', B'\}$, $\{\{x\}, X\}$ together contradict the profile property. □

We can also deduce statements analogue to Theorem 7 and Theorem 8 for profiles, i.e., we can consider regular profiles on the set $\vec{S}(E)$ of separations of the edges of our bipartite graph and show the following:

Theorem 11. *If P is a regular profile of $\vec{S}_{2k}(E)$, then $P_{X,k}$ is a regular profile of $\vec{S}_k(X)$.*

Proof. If $k \leqslant \frac{1}{2}$, then P_X only orients separations (A, B) of order less than $\frac{1}{2}$, that is, only separations where each vertex in Y has all its neighbours in either $A \setminus B$ or $B \setminus A$. It is then easy to see that P_X is indeed a profile.

So, suppose that $k > \frac{1}{2}$. We first note that the set $P_{X,k}$ contains at least one of (A, B) and (B, A) for every separation $(A, B) \in \vec{S}_k(X)$. Indeed, by Proposition 4.9 $|(A, B)^E|_E \leqslant 2|A, B|_X$, and so since P is an orientation of $\vec{S}_{2k}(E)$ either $(A, B)^E \in P$ or $(B, A)^E \in P$.

Let us now show that for no separation $\{ A, B \}$ we have both (A, B) and (B, A) in $P_{X,k}$. Suppose otherwise, then P contains separations (C_1, D_1) and (C_2, D_2) such that $(C_1, D_1)^{\blacktriangleleft} = (A, B)$ and $(C_2, D_2)^{\blacktriangleleft} = (B, A)$.

Note that, by Proposition 4.9, we have $|E(A), E(B)|_E \leqslant 2|A, B|_X < 2k$, and, hence, $(A, B)^E \in P$ or $(B, A)^E \in P$. As $(E(A), E(B))^{\blacktriangleleft} = (A, B)$, we may suppose without loss of generality that either $(C_1, D_1) = (A, B)^E$ or $(C_2, D_2) = (B, A)^E$. We suppose the former one, the latter case is similar.

Now pick a separation $(C, D) \in P$ so that $(C, D)^{\blacktriangleleft} \geqslant (C_2, D_2)^{\blacktriangleleft} = (B, A)$ and the set $(D_1 \cap D) \setminus (C_1 \cup C)$ is as small as possible. We claim that $C_1 \cup C \neq E$. So suppose for a contradiction that $C_1 \cup C = E$. Then $(C_1, D_1) \vee (C, D) = (C_1 \cup C, D_1 \cap D) = (E, D_1 \cap D)$. Since $D_1 \subseteq E(B)$ and $D \subseteq E(A)$ we have that $D_1 \cap D \subseteq E(A \cap B)$. We thus get

$$|E, D_1 \cap D|_E \leqslant |D_1 \cap D| \leqslant |E(A \cap B)| \leqslant 2|A, B|_X < 2k,$$

which gives a contradiction, since $(E, D_1 \cap D) \in P$ contradicts the fact that P is regular. Thus, $C_1 \cup C \neq E$ and there exists some edge $e \in (D_1 \cap D) \setminus (C_1 \cup C)$.

Let x be the end vertex of e in X. Note that $e \in E(B) \setminus E(A)$ since $E(A) = C_1$ and $e \notin C_1$. Thus, $x \in B \setminus A$. Moreover, as $B = C_2{}^{\blacktriangleleft}_{D_2} \subseteq C^{\blacktriangleleft}_D$, we have that $x \in C^{\blacktriangleleft}_D$. Thus, e is incident with $C^{\blacktriangleleft}_D$.

Consequently, we can apply Lemma 4.13 to get that $|C + e, D| \leqslant |C, D|$. Thus, P orients $(C + e, D)$ and therefore $(C + e, D) \in P$ as $(D, C + e) \in P$ would contradict the profile property, since $(C, D) \vee (\{e\}, E) = (C + e, D)$ and $|\{e\}, E|_X \leqslant 1$ and consequently $(\{e\}, E) \in P$, since P is regular and $k > \frac{1}{2}$.

But this implies that $(C + e, D) \in P$ is a better choice for (C, D), as $(C + e, D)^{\blacktriangleleft} \geqslant (C, D)^{\blacktriangleleft}$ by Lemma 4.7 and $(D_1 \cap D) \setminus (C_1 \cup C) \supsetneq (D_1 \cap D) \setminus (C_1 \cup (C + e))$ as $e \in (D_1 \cap D) \setminus (C_1 \cup C)$.

Thus, $P_{X,k}$ is indeed an orientation. It is also consistent, for if $(A_1, B_1) \leqslant (A_2, B_2)$, then $(A_1, B_1)^E \leqslant (A_2, B_2)^E$ and thus, if $(A_2, B_2) \in P_X \cap \vec{S}_k(X)$ and $(A_1, B_1) \in \vec{S}_k(X)$ then $(A_1, B_1)^E \in P$ by Proposition 4.9 and thus $(A_1, B_1) \in P_X$. In particular, $(B_1, A_1) \notin P_{X,k}$, as $P_{X,k}$ is a orientation.

That $P_{X,k}$ is a profile now follows like this: suppose that (A_1, B_1), (A_2, B_2), $(B_1 \cap B_2, A_1 \cup A_2)$ is a triple in $P_{X,k}$ contradicting the profile property, then P

orients $(A_1, B_1)^E$, $(A_2, B_2)^E$ and $(B_1 \cap B_2, A_1 \cup A_2)^E$ by Proposition 4.9. By the above observation, P orients them as $(A_1, B_1)^E$,$(A_2, B_2)^E$ and $(B_1 \cap B_2, A_1 \cup A_2)^E$, since $P_{X,k}$ does not contain any $(((A_i, B_i)^E)^*)^{\blacktriangleleft} = (B_i, A_i)$. However, the three separations $(A_1, B_1)^E$, $(A_2, B_2)^E$ and $(B_1 \cap B_2, A_1 \cup A_2)^E = (B_1, A_1)^E \wedge (B_2, A_2)^E$ in P then contradict the profile property.

The profile $P_{X,k}$ is regular, since if $(X, A) \in P_{X,k}$, then $(E(X), E(A)) = (E, E(A))$ is a cosmall separation in P, which contradicts the regularity of P. □

Theorem 12. *If P is a regular profile of $\vec{S}_{3k}(X)$, then $P_{E,k}$ is a regular profile of $\vec{S}_k(E)$.*

Proof. By Lemma 4.10, given some separation $(C, D) \in \vec{S}_k(E)$ we have that $|(C, D)^{\blacktriangleleft}|_X \leqslant |C, D|_E$, thus P contains exactly one of $(C, D)^{\blacktriangleleft}$ and $((C, D)^{\blacktriangleleft})^* = (D, C)^{\blacktriangleleft}$, and consequently $P_{E,k}$ contains exactly one of (C, D) and (D, C), i.e., $P_{E,k}$ is an orientation of $\vec{S}_k(E)$. That the orientation $P_{E,k}$ is consistent is then immediate from Lemma 4.7.

So, it remains to show that $P_{E,k}$ satisfies the profile property. Let us suppose for a contradiction that there is some set

$$\{\,(C_1, D_1), (C_2, D_2), (C_3, D_3)\,\} \subseteq P_{E,k}$$

such that $(C_1, D_1) \vee (C_2, D_2) = (D_3, C_2)$. Let $(A_i, B_i) = (C_i, D_i)^{\blacktriangleleft}$ for each $i = 1, 2, 3$ and let $Z = B_1 \cap B_2 \cap B_3$.

Since $Z \subseteq B_i = D_i{}^{\blacktriangleleft}_{C_i}$ for each i, we have that $|E(z) \cap D_i| \geqslant |E(z) \cap C_i|$ for all $z \in Z$ and $i = 1, 2, 3$. However, since $C_1 \cup C_2 \cup C_3 = E$,

$$\begin{aligned}
\sum_{i=1}^{3} |C_i, D_i|_E &= \sum_{i=1}^{3} \sum_{v \in V} \Big(\mathsf{min}\{\, |E(v) \cap C_i|\,, |E(v) \cap D_i|\,\} - \frac{1}{2}\,|E(v) \cap C_i \cap D_i| \Big) \\
&\geqslant \sum_{i=1}^{3} \sum_{z \in Z} \Big(\mathsf{min}\{\, |E(z) \cap C_i|\,, |E(z) \cap D_i|\,\} - \frac{1}{2}\,|E(z) \cap C_i \cap D_i| \Big) \\
&= \sum_{z \in Z} \sum_{i=1}^{3} \Big(|E(z) \cap C_i| - \frac{1}{2}\,|E(z) \cap C_i \cap D_i| \Big) \\
&\geqslant \sum_{z \in Z} d(z)/2 = \frac{|E(Z, Y)|}{2}.
\end{aligned}$$

As $|C_i, D_i|_E < k$ for every $i = 1, 2, 3$, this gives us $|E(Z, X)| < 6k$ and thus

$$|Z, X|_X = \frac{|E(Z, Y)|}{2} < 3k.$$

Hence, P needs to contain (Z, X), as P is a regular $3k$-profile.

Moreover, P contains $(A_1, B_1) \vee (A_2, B_2)$ as by submodularity

$$|(A_1, B_1) \vee (A_2, B_2)|_X \leqslant |A_1, B_1|_X + |A_2, B_2|_X < 2k.$$

Then also $(A_1, B_1) \vee (A_2, B_2) \vee (A_3, B_3) \in P$ as

$$|(A_1, B_1) \vee (A_2, B_2) \vee (A_3, B_3)|_X \leqslant |(A_1, B_1) \vee (A_2, B_2)|_X + |A_3, B_3|_X < 3k.$$

But $(A_1, B_1) \vee (A_2, B_2) \vee (A_3, B_3) = (A_1 \cup A_2 \cup A_3, Z)$ and, as $Z = B_1 \cap B_2 \cap B_3$ we have that $A_1 \cup A_2 \cup A_3 \cup Z = X$ and thus $(A_1 \cup A_2 \cup A_3, Z) \vee (Z, X) = (X, Z)$ which contradicts the fact that P is a profile.

The profile $P_{E,k}$ is regular, since if $(E, C) \in P_{E,k}$, then $(E, C)^{\blacktriangleright} = (X, C_E^{\blacktriangleright})$ is a cosmall separation in P, which contradicts the regularity of P. □

Another possible variation of the problem is to consider other ways to relate tangles of the different systems to each other. Given our shifting operation between the two separation systems $\vec{S}(X)$ and $\vec{S}(Y)$ we defined a 'pull-back' type operation that maps subsets of $\vec{S}(X)$ to subsets of $\vec{S}(Y)$ and investigated its action on tangles. However, as in the definition of τ_X there is another way to extend our shifting operations from acting on single separations to acting on subsets via a 'push-forward' type action. It is perhaps equally natural to ask how the tangles of $\vec{S}_k(X)$ and $\vec{S}_k(Y)$ behave under these operations.

Given a tangle τ of $\vec{S}_k(X)$ one may define the set

$$\tau^{\triangleright} := \{\, (A, B)^{\triangleright} : (A, B) \in \tau \,\} \subseteq \vec{S}_k(Y),$$

and similarly, if τ is a tangle of $\vec{S}_k(Y)$, we may define

$$\tau^{\triangleleft} := \{\, (A, B)^{\triangleleft} : (A, B) \in \tau \,\} \subseteq \vec{S}_k(X).$$

Note that $\tau^{\triangleright}$ and $\tau^{\triangleleft}$, generally, are no more than subsets of $\vec{S}_k(Y)$ or $\vec{S}_k(X)$, respectively, they need not be an orientation, not even a partial orientation.

However, we can show that this push-forward $\tau^{\triangleright}$ is, when restricted appropriately, contained in a corresponding pull-back ${}^{\triangleright}\tau$ and thus needs to be a partial orientation satisfying (†).

Proposition 4.17. *Let τ be a tangle of $\vec{S}_{16k}(X)$, then $(\tau \cap \vec{S}_k(X))^{\triangleright} \subseteq {}^{\triangleright}\tau$.*

Proof. The only way in which this may fail is that for some $(C, D) \in {}^{\triangleright}\tau$ we have $(D, C) \in (\tau \cap \vec{S}_k(X))^{\triangleright}$. Let us say this happens because of some $(A, B) \in \tau \cap \vec{S}_k(X)$ with $(A, B)^{\triangleright} = (D, C)$.

Then also $(A, B) \in {}^{\triangleleft}_{k}({}^{\triangleright}_{4k}\tau)$ by Theorem 6, and hence $(A, B)^{\triangleleft} = (D, C) \in {}^{\triangleright}_{4k}\tau$, contradicting the fact that ${}^{\triangleright}_{4k}\tau$ is a tangle. □

A third variation of this idea comes from applications. There we often wish to work with systems of bipartitions, rather than more general set separations. Again here much of the work in previous sections remains true in this setting, with slight tweaks to the definitions and results.

More explicitly, given as before a bipartite graph G on bipartition classes X and Y let $\vec{\mathcal{B}}(X)$ and $\vec{\mathcal{B}}(Y)$ be the universe of all bipartition of X and Y respectively.

Given a bipartition (A, B) of X, we can define, as before, the shift of (A, B) to be the bipartition (C, D) of Y where C is the set of all elements of Y with more neighbours in A than in B and D is the set of all elements of Y with more neighbours in B than in A. However, a small issue arises here as to what to do with those vertices which have an equal number of neighbours in A and B. Since we need the shift of a bipartition to be a bipartition we need to break the symmetry in some way here. We define our shifting operation not for unoriented, but for oriented bipartitions, namely we define a bipartition $(A, B)^{\triangleright} := (C, D)$ of Y by letting

$$C := \{\, y \in Y : |N(y) \cap A| \geqslant |N(y) \cap B| \,\}$$

and

$$D := \{\, y \in Y : |N(y) \cap A| < |N(y) \cap B| \,\}.$$

In particular, in general $(A, B)^{\triangleright} \neq ((B, A)^{\triangleright})^*$.

There is again a natural order function for these bipartitions given by

$$|A, B|_X := \sum_{y \in Y} \mathsf{min}\{\, |N(y) \cap A| \,, |N(y) \cap B| \,\},$$

which can again be seen to be submodular,[3] and for a suitable definition of a tangle we can show that analogues of Theorem 5 and 6 hold for tangles of $\vec{\mathcal{B}}(X)$ and $\vec{\mathcal{B}}(Y)$. In particular the following is true:

Theorem 4.18. *Let τ be a tangle of $\vec{\mathcal{B}}_{4k}(X)$, then $\tau' := {}^{\triangleright}_{k}\tau$ is a tangle of $\vec{\mathcal{B}}_k(Y)$.*

Theorem 4.19. *Let τ be a tangle of $\vec{\mathcal{B}}_{16k}(X)$, let $\tau' = {}^{\triangleright}_{4k}\tau$, and let $\tau'' = {}^{\triangleleft}_{k}\tau'$, then $\tau'' \subseteq \tau$.*

Again, the arguments closely follow the proofs of Theorem 5 and 6, for details see the extended version of the paper [24] upon which this chapter is based.

It would be nice if one could find a unified result implying these different variations. Unfortunately, it seems that the nature of the result means that strengthening or weakening the notion of tangle we consider does not make the statement stronger or weaker, but rather incomparable. Indeed, since we wish to show that tangles of $\vec{S}(X)$ shift to tangles of $\vec{S}(Y)$, if we consider a stronger notion of tangle, then fewer orientations are tangles, and so it is required to show that a stronger property holds for the shifts, but under a stronger assumption on the original orientations. Similarly, if we consider a weaker notion of tangles, then more orientations will be tangles, and so while it is required to show that only the weaker property holds for the shifts, we also only have weaker assumptions on the original orientations.

A similar problem arises if one wants to relate the statements for bipartitions to the statements about set separations: In principle, every tangle of set separations induces a tangle of bipartitions. Conversely, every tangle of bipartitions induces a tangle of set separations of lower order, except that the ‘regularity’ conditions of these two types of tangles are not compatible: For set separations we just require that we do not contain any cosmall separations, whereas for bipartitions we want more, namely that the big side of our bipartition of the edges meet both sides of the graph in at least two vertices. Thus, the statements for these two types of tangles are formally independent, although most of the proof strategy is very similar.

Part II.

Trees of tangles

5. Trees of tangles and the splinter lemmas

In this chapter we will present a unified approach to proving tree-of-tangles theorems. Sections 5.1 to 5.3 are based on [38], which is joint work with Christian Elbracht and Jakob Kneip. Sections 5.6 and 5.7 is as-of-yet unpublished collaboration with Christian Elbracht. The extension to the infinite in Sections 5.8 and 5.9 is based on [41], which also is joint work with Christian Elbracht and Jakob Kneip. The application to edge blocks, Section 5.10, is based on the preprint [42] which is joint work with Christian Elbracht and Jan Kurkofka.

5.1. Trees of tangles in abstract separation systems

The central theorem in the theory of tangles, which was established by Robertson and Seymour together with the notion of tangles, is the following:

Theorem 5.1 ([64]). *Every finite graph has a tree-decomposition displaying its maximal tangles.*

Theorem 5.1 roughly says that the highly cohesive regions in a graph are arranged in a tree-like structure. The 'maximal' in Theorem 5.1 relates to the order of the tangles: the tree-decomposition found by Theorem 5.1 displays the graph's tangles at every level of coarseness.

The original proof of Theorem 5.1 by Robertson and Seymour in [64] is fairly involved and uses as tools multiple non-trivial results about separations in graphs, for instance, the existence of certain 'tie-breaker' functions. Since then, the theory of tangles has moved on considerably, and shorter and more elementary proofs of Theorem 5.1 have been found. The shortest proof to date is due to Carmesin [9, 19], who utilizes the fact that the separations needed for the tree-decomposition in Theorem 5.1 behave well under joins and meets when taken to be of minimal order.

Carmesin, Diestel, Hundertmark, and Stein established the following variant on Theorem 5.1:

Theorem 5.2 ([13]). *For every integer $k \geqslant 0$, every finite graph has a canonical tree-decomposition displaying its k-blocks.*

While 'k-blocks' are a significant sidestep from 'maximal tangles', the notable part is the caonicity. Here, 'canonical' means that every automorphism of the graph acts on the decomposition tree. In other words, Theorem 5.2 uses only invariants of the graph – in particular no tie-breaker – to find the desired tree-decomposition.

A careful analysis of the proof of Theorem 5.2, started in [11, 12] resulted in another shift of paradigm in [26], similarly to the shift brought about by [64]. Much in the same way as tangles made it possible indirectly to capture substructures in graphs that were traditionally described more directly by sets of vertices or edges, and to treat them in a unified framework, it turned out that tangles themselves could be described, unlike in their definition given by Robertson and Seymour in [64], without reference to vertices or edges. They prove:

Theorem 5.3 ([26])**.** *Every finite graph has a canonical tree-decomposition displaying its robust regular profiles.*

Indeed, the only information needed about a graph's tangles to prove Theorem 5.3 is how its separations relate to each other, that is, which separations are nested or cross. This information, a partial order between the separations of the graph together with an involution became the definition of an abstract separation system, and all subsequent tools and theorems in [26] are then formulated for separation systems. This new abstract representation of tangles yielded not only a cleaner proof of Theorem 5.3 in [26], but also made the theory of tangles applicable to a wider range of combinatorial structures.

However one condition not expressed in terms of relations between the separations remained in use throughout the series of abstractions of Theorem 5.1 implemented in [13] and [26]: all of these works assumed that the separation systems of interest came with a submodular order function. Likewise, Carmesin's short proof of Theorem 5.1 in [9] also leverages the fact that the order of separations of graphs is a submodular function.

This last non-structural aspect of tree-of-tangles theorems was disposed of in [25]: in that paper Diestel, Erde, and Weißauer replaced the order function with a purely structural notion of submodularity which can be expressed solely in terms of the lattice structure of a universe of separation surrounding the separation system. In doing so they established the most general and widely applicable variant of Theorem 5.1 to date:

Theorem 5.4 ([25, Theorem 6])**.** *Let $\vec{S}$ be a submodular separation system in some universe $\vec{U}$ of separations and let $\mathcal{P}$ a set of profiles of S. Then S contains a nested set that distinguishes $\mathcal{P}$.*

Since the relevant separation systems in graphs are all 'structurally submodular', Theorem 5.4 still applies to tangles in graphs. On the other hand there are separation systems that are structurally submodular but cannot be represented by graph separations ([4]) or by a submodular order function (cf. Section 8.3). In particular, Theorem 5.4 can also be applied to separation systems which, unlike separations in graphs, do not come with any order function, such as arbitrary bipartitions of sets. This is a marked step forward from its predecessor Theorem 5.1, whose original proof made heavy use of the order of particular separations.

However there is a trade-off involved in Theorem 5.4's wider applicability: it does not imply Theorem 5.1. Indeed, Theorem 5.4 applied to a graph produces a tree-decomposition which displays just the graph's k-tangles for arbitrary but fixed k. This is a significant weakening of Theorem 5.1, which finds a decomposition displaying the graph's maximal tangles for all tangles orders simultaneously. Moreover, the tree-decomposition found by Theorem 5.1 is *efficient* in the sense that for every pair of tangles distinguished by the tree-decomposition, the separation in the decomposition distinguishing that pair of tangles is of the lowest possible order. Since Theorem 5.4 makes only structural assumptions so as to be applicable to separation systems without any order function, Theorem 5.4 cannot guarantee that the separations used by the nested set to distinguish a particular pair of tangles are of minimal order.

In this chapter we bridge the gap between Theorem 5.1 and Theorem 5.4 by establishing the following tree-of-tangles theorem which combines the upsides of both Theorem 5.1 and Theorem 5.4, i.e., which is as widely applicable as Theorem 5.4 while still being as powerful and efficient as Theorem 5.1 when applied to tangles in graphs:

Theorem 14. *If $\mathcal{S} = (S_1, \dots, S_n)$ is a compatible sequence of structurally submodular separation systems inside a universe U, and $\mathcal{P}$ is a robust set of profiles in $\mathcal{S}$, then there is a nested set N of separations in U which efficiently distinguishes all the distinguishable profiles in $\mathcal{P}$.* ⇝ p. 72

Theorem 14 includes Theorem 5.4 by taking a sequence of just one separation system, and it implies Theorem 5.1 by taking as separation systems S_k the sets of all separations of order $< k$ of the given graph; the resulting nested set is the set of separations of the desired tree-decomposition.

The nested set N found by Theorem 14 has to contain for every pair of profiles in $\mathcal{P}$ a separation from that pair's 'candidate set' of all those separations which (efficiently) distinguish that pair of profiles. Thus, to prove Theorem 14, it suffices to show that one can pick an element from each of these 'candidate sets' in a nested way.

As it turns out, there is a very simple and purely structural requirement of the way these 'candidate sets' interact with each other which guarantees that it is possible to pick such a nested set:

Lemma 13 (Splinter Lemma). *Let U be a universe of separations and $\mathfrak{A} = (\mathcal{A}_i)_{i \leqslant n}$ a family of subsets of U. If $\mathfrak{A}$ splinters, then we can pick an element a_i from each $\mathcal{A}_i$ so that $\{ a_1, \dots, a_n \}$ is nested.* ⇝ p. 64

Lemma 13, in a sense, represents yet another step of abstraction in the theory of tangles: rather than working with the profiles themselves it works with the sets of separations distinguishing a given pair of profiles.

Lemma 13 not only implies Theorem 14, but can also be used to prove Theorem 5.1 and Theorem 5.4 directly. In fact Lemma 13 has a remarkably short proof (as we shall see in Section 5.2), making it the shortest available proof of Theorem 5.1 so far (see Section 5.3.1). Moreover, the premise in Lemma 13 is straightforward to check, and Lemma 13 itself does not make reference to tangles or any specific implementations of them. As a result Lemma 13 can be used in many different settings, implying variations of Theorem 5.1 in a multitude of contexts. For example, after deriving in Section 5.3 Theorem 5.1, Theorem 5.4, and Theorem 14 from Lemma 13, we use Lemma 13 to establish a new tree-of-tangles theorem in the setting of clique separations.

Since Lemma 13 does not yield a canonical set of separations, we cannot deduce Theorem 5.3 from Lemma 13. We fix this in Section 5.4 by establishing a version of Lemma 13 which does give a canonical nested set, albeit under slightly stronger assumptions:

⇝ p. 77 **Lemma 16** (Canonical splinter lemma)**.** *Let U be a universe of separations and let $\mathfrak{A} = (\mathcal{A}_i : i \in I)$ be a collection of subsets of U that splinters hierarchically with respect to a partial order $\preccurlyeq$ on I. Then there exists a nested set $N = N(\mathfrak{A})$ meeting every $\mathcal{A}_i$ in $\mathfrak{A}$.*

Moreover, $N(\mathfrak{A})$ is canonical: if φ is an isomorphism of separation systems between $\bigcup_{i\in I} \overrightarrow{\mathcal{A}}_i$ and a subset of some universe U' such that the family $\varphi(\mathfrak{A}) := (\varphi(\mathcal{A}_i) : i \in I)$ splinters hierarchically with respect to $\preccurlyeq$, then $N(\varphi(\mathfrak{A})) = \varphi(N(\mathfrak{A}))$.

We make use of Lemma 16 in Section 5.5 to obtain a new shortest proof of Theorem 5.3 and to extend Theorem 5.3 to two natural types of separations whose structural submodularity does not come from a submodular order function: clique separations, and circle separations. Only very recently, a theorem was established by Elbracht and Kneip which, so far, could not be shown using this method: a canonical version of Theorem 5.4 [34]

We will later, in Section 5.6, uncover that in its essence Lemma 13 is just a lemma about the nestedness relation and can be applied to any similarly behaved relation. We will demonstrate how this can be applied to directed tangles in directed graphs.

Lemma 16 can be interpreted as a lemma about a relation in the same way, and that is the approach we will take when establishing the 'thin splinter lemma' for infinite settings in Section 5.8.

5.2. The Splinter Lemma

In this section we establish our first main result, Lemma 13, from which we shall derive two previously known theorems as well as two new flavours of tree-of-tangles theorems in Section 5.3. A cornerstone of the proofs of both Lemma 13 as well as of the two known results we shall derive from it is the so-called 'fish lemma':

Lemma 2.1 ('Fish lemma', [20, Lemma 3.2]). *Let U be a universe of separations and $r, s \in U$ be two crossing separations. Every separation $t \in U$ that is nested with both r and s is also nested with all four corner separations of r and s.* ⇝ p. 13

Typically, the proof of a tree-of-tangles theorem proceeds by starting with some set N of separations which distinguish some (or all) of the given tangles, and then repeatedly replacing elements r of N which cross some other element s of N with an appropriate corner separation of r and s. Lemma 2.1 is then used to show that each of these replacements makes N 'more nested', and thus one eventually obtains a nested set N which distinguishes all the given tangles. (See for instance the proof of Theorem 4 of [25].) Usually, in order to not reduce the set of tangles distinguished by N, one has to take special care which corner separation of two crossing r and s in N one uses for replacement; this depends on the specific properties of the tangles at hand.

Our Lemma 13 seeks to eliminate this careful selection of corner separations for replacement: we will show that for a family $(\mathcal{A}_i)_{i \leqslant n}$ of subsets of some universe U we can find a nested set N meeting all the $\mathcal{A}_i$, provided that these sets $\mathcal{A}_i$ have one straightforward-to-check property. This lemma will imply many of the existing tree-of-tangles theorems by taking as sets $\mathcal{A}_i$ the sets of separations which distinguish the i-th pair of tangles, and checking that the one assumption needed for Lemma 13 is met. Notably, Lemma 13 will make no reference at all to tangles or their specific properties. The proof of Lemma 13 will also utilize Lemma 2.1; however, the only assumption we need about the sets $\mathcal{A}_i$ is that for elements a_i and a_j of $\mathcal{A}_i$ and $\mathcal{A}_j$, respectively, one of their four corner separations lies in either $\mathcal{A}_i$ or $\mathcal{A}_j$. This condition will be easy to verify if one wants to deduce other tree-of-tangles theorems from Lemma 13. In fact, the verification of this condition, which just asks for the existence of *some* corner separation of a_i and a_j in $\mathcal{A}_i \cup \mathcal{A}_j$, will usually be much more straightforward than the hands-on arguments used in the original proofs of those tree-of-tangles theorems, which for their replacement arguments often need to prove the existence of a *specific* corner separation of a_i and a_j. So let us define this condition formally.

Let U be a universe and $\mathfrak{A} = (\mathcal{A}_i)_{i \leqslant n}$ some family of non-empty subsets of U. We say that $\mathfrak{A}$ *splinters* if, for every crossing pair of $a_i \in \mathcal{A}_i \setminus \mathcal{A}_j$ and $a_j \in \mathcal{A}_j \setminus \mathcal{A}_i$, one of their four corner separations lies in $\mathcal{A}_i \cup \mathcal{A}_j$.

Observe that a family $(\mathcal{A}_i)_{i \leqslant n}$ of non-empty sets splinters if and only if for every pair $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$ of separations, either some corner separation of a_i and a_j lies in $\mathcal{A}_i \cup \mathcal{A}_j$, or one of a_i and a_j lies in $\mathcal{A}_i \cap \mathcal{A}_j$. This is, because if two separations a_i and a_j are nested, then these separations themselves are corner separations of the pair a_i and a_j.

With this definition and Lemma 2.1 we are already able to state and prove our first main result:

Lemma 13 (Splinter Lemma). *Let U be a universe of separations and $\mathfrak{A} = (\mathcal{A}_i)_{i \leqslant n}$ a family of subsets of U. If $\mathfrak{A}$ splinters, then we can pick an element a_i from each $\mathcal{A}_i$ so that $\{ a_1, \dots, a_n \}$ is nested.*

Proof. We proceed by induction on n. The assertion clearly holds for $n = 1$. So suppose that $n > 1$ and that the above assertion holds for all smaller values of n.

Suppose first that we can find some $a_i \in \mathcal{A}_i$ so that a_i is nested with at least one element of $\mathcal{A}_j$ for each $j \neq i$. Then the assertion holds: for $j \neq i$ let $\mathcal{A}'_j$ be the set of those elements of $\mathcal{A}_j$ that are nested with a_i. Then $(\mathcal{A}'_j : j \neq i)$ is a family of non-empty sets which splinters by Lemma 2.1. Thus by the induction hypothesis we can pick a nested set $\{ a_j \in \mathcal{A}'_j : j \neq i \}$, which together with a_i is the desired nested set.

To conclude the proof it thus suffices to find an a_i as above. To this end, we apply the induction hypothesis to $\mathcal{A}_1, \dots, \mathcal{A}_{n-1}$ to obtain a nested set consisting of some $a_1, \dots, a_{n-1}$. Fix an arbitrary $a_n \in \mathcal{A}_n$. For all $i < n$, if a_i itself or one of its corner separations with a_n lies in $\mathcal{A}_n$, this a_i is the desired separation for the above argument. Otherwise, for each $i < n$, either a_n itself or one of its corner separations with a_i lies in $\mathcal{A}_i$, in which case a_n is the desired separation for the above argument. $\square$

We shall see in Section 5.3 that this innocuous-looking lemma is actually strong enough to directly imply various existing tree-of-tangles theorems, including Theorem 5.1.

Although Lemma 13 is already strong enough to imply various tree-of-tangles theorems – as we will show in Section 5.3 – we can prove an even stronger version of Lemma 13 with no additional assumptions:

Lemma 5.5. *Let U be a universe, $\mathfrak{A} = (\mathcal{A}_i)_{i \leqslant n}$ a family of subsets of U, and x_i some fixed element of $\mathcal{A}_i$ for each $i \leqslant n$. If $\mathfrak{A}$ splinters then we can pick an element a_i from each $\mathcal{A}_i$ so that $\{ a_1, \dots, a_n \}$ is nested, with the additional property that every $s \in U$ which is nested with all x_i is also nested with all a_i.*

Proof. We proceed by induction on n. The assertion clearly holds for $n = 1$ by picking x_1 as a_1.

So suppose that $n > 1$ and that the above assertion holds for $n - 1$. Let $x_1, \dots, x_n$ be fixed elements from $\mathcal{A}_1, \dots, \mathcal{A}_n$ respectively. By applying the induction hypothesis to $\mathcal{A}_1, \dots, \mathcal{A}_{n-1}$ and $x_1, \dots, x_{n-1}$, we can pick elements $a_i \in \mathcal{A}_i$ for $i < n$ so that $\{ a_1, \dots, a_{n-1} \}$ is nested and every separation of U which is nested with $\{ x_1, \dots, x_{n-1} \}$ is also nested with $\{ a_1, \dots, a_{n-1} \}$.

Among all separations in $\mathcal{A}_n$ which are nested with each separation of U that is nested with all of $\{ x_1, \dots, x_n \}$, pick a separation y_n which crosses a_i for as few $i < n$ as possible. Note that x_n is a candidate for this y_n.

If y_n is nested with $\{ a_1, \dots, a_{n-1} \}$, we are done, as then $\{ a_1, \dots, a_{n-1}, y_n \}$ is as desired. So suppose that, after re-arranging, y_n crosses $a_1, \dots, a_k$ for some $k < n$ and is nested with $a_{k+1}, \dots, a_{n-1}$. Then, for each $i \leqslant k$, neither a_i itself nor any of its corner separations with y_n can lie in $\mathcal{A}_n$: they would all have been better choices for y_n by Lemma 2.1. Therefore, for every $i \leqslant k$, either y_n itself or one of its corner separations with a_i lies in $\mathcal{A}_i$. Let us write a_i' for that separation.

Observe that y_n is nested with $\{ a_1', \dots, a_k', a_k, \dots, a_{n-1} \}$, and so is every separation of U that is nested with all of $\{ x_1, \dots, x_n \}$ by Lemma 2.1. We are not done yet, however, since $\{ a_1', \dots, a_k' \}$ might not be a nested set.

To finish the proof, we apply the induction hypothesis to $\mathcal{A}_1, \dots, \mathcal{A}_{n-1}$ a second time, this time using $a_1', \dots, a_k', a_{k+1}, \dots, a_{n-1}$ as the fixed elements. In doing so we obtain a choice of $b_i \in \mathcal{A}_i$ for every $i < n$ so that $\{ b_1, \dots, b_{n-1} \}$ is a nested set that is also nested with each separation of U which is nested with $\{ a_1' \dots, a_k', a_{k+1}, \dots, a_{n-1} \}$. We claim that $\{ b_1 \dots, b_{n-1}, y_n \}$ is the desired set to finish the induction step.

The separation y_n is nested with $\{ b_1, \dots, b_{n-1} \}$ since it is nested with the fixed elements $a_1', \dots, a_k', a_{k+1}, \dots, a_{n-1}$. Furthermore, let $s \in U$ be a separation that is nested with $x_1, \dots, x_n$. Then s is nested with y_n by choice of y_n as well as with $a_1 \dots, a_{n-1}$ by the induction hypothesis. Hence s is nested with $a_1', \dots, a_k', a_{k+1}, \dots, a_{n-1}$ by Lemma 2.1, and therefore with $\{ b_1, \dots, b_{n-1} \}$. $\square$

5.3. Applications of the Splinter Lemma

5.3.1. A short proof of Theorem 5.1

As a first application of Lemma 13 let us give a short proof of Theorem 5.1:

⇝ p. 59 **Theorem 5.1** ([64]). *Every finite graph has a tree-decomposition displaying its maximal tangles.*

We say that a tree-decomposition $(T, \mathcal{V})$ of a graph G *displays its maximal tangles* if the set of separations induced by $(T, \mathcal{V})$ efficiently distinguishes the set of all maximal tangles of G.

If N is a nested set of separations of G it is straightforward to find a tree-decomposition of G whose set of induced separations is precisely N (see [21, 64]). Therefore, in order to prove Theorem 5.1, it suffices to find a nested set N of separations of G which efficiently distinguishes all maximal tangles of G.

For every pair τ, τ' of distinct maximal tangles of G let

$$\mathcal{A}_{\tau,\tau'} := \{\, \{A, B\} \in S(G) : \{A, B\} \text{ efficiently distinguishes } \tau \text{ and } \tau' \,\}.$$

Since P and P' are not subsets of each other $\mathcal{A}_{\tau,\tau'}$ is a non-empty set.

Let $\mathfrak{A}$ be the family of all these sets $\mathcal{A}_{\tau,\tau'}$. A nested set of separations of G distinguishes all maximal tangles of G efficiently if and only if it contains an element of each $\mathcal{A}_{\tau,\tau'}$. Therefore the existence of such a set, and hence Theorem 5.1, now follows directly from Lemma 13 once we show that $\mathfrak{A}$ splinters:

Lemma 5.6. *The family $\mathfrak{A}$ of all $\mathcal{A}_{\tau,\tau'}$ splinters.*

Proof. Let $\tau \neq \tau'$ and $\sigma \neq \sigma'$ be two pairs of distinct maximal tangles of G and let $\{\, A, B \,\} \in \mathcal{A}_{\tau,\tau'}$ and $\{\, C, D \,\} \in \mathcal{A}_{\sigma,\sigma'}$ be two crossing separations. We need to show that we have either $\{\, A, B \,\} \in \mathcal{A}_{\sigma,\sigma'}$ or $\{\, C, D \,\} \in \mathcal{A}_{\tau,\tau'}$, or that some corner separation of $\{\, A, B \,\}$ and $\{\, C, D \,\}$ lies in $\mathcal{A}_{\tau,\tau'} \cup \mathcal{A}_{\sigma,\sigma'}$. By switching their roles if necessary we may assume that $|(A, B)| \leqslant |(C, D)|$.

Since σ and σ' both orient (C, D), and $|(A, B)| \leqslant |(C, D)|$, both tangles also orient $\{\, A, B \,\}$. If σ and σ' orient $\{\, A, B \,\}$ differently, then $\{\, A, B \,\}$ distinguishes them efficiently and hence lies in $\mathcal{A}_{\sigma,\sigma'}$. So suppose that σ and σ' contain the same orientation of $\{\, A, B \,\}$, say, (A, B).

By renaming them if necessary, we may assume that $(C, D) \in \sigma$ and $(D, C) \in \sigma'$.

For the corner separation $(A \cup C\,,\, B \cap D)$, we first consider the case that $|(A \cup C\,,\, B \cap D)| \leqslant |(C, D)|$. Then, by $(A, B), (C, D) \in \sigma$ and the tangle property (T), Q must contain $(A \cup C\,,\, B \cap D)$. On the other hand σ' must contain its inverse $(B \cap D\,,\, A \cup C)$ since $(D, C) \in Q'$. But then this corner separation efficiently distinguishes σ and σ' and hence lies in $\mathcal{A}_{\sigma,\sigma'}$.

Thus we may suppose that $|(A \cup C\,,\, B \cap D)| \geqslant |(C, D)|$. By a similar argument we may further suppose that $|(A \cup D\,,\, B \cap C)| \geqslant |(C, D)|$. Submodularity then yields $|(A \cap C\,,\, B \cup D)|\,, |(A \cap D\,,\, B \cup C)| \leqslant |(A, B)|$.

By switching the roles of τ and τ' if necessary, we may assume that $(A, B) \in \tau$ and $(B, A) \in \tau'$. Then, by the above inequality, τ must contain both $(A \cap C\,,\ B \cup D)$ and $(A \cap D\,,\ B \cup C)$, since it cannot contain either of their inverses due to $(A, B) \in \tau$ and the tangle property (T). However, due to $(B, A) \in \tau'$ and the tangle property (T), τ' cannot contain both of $(A \cap C\,,\ B \cup D)$ and $(A \cap D\,,\ B \cup C)$. In must therefore contain the inverse of at least one of these corner separations, which then efficiently distinguishes τ and τ' and hence lies in $\mathcal{A}_{\tau,\tau'}$. □

5.3.2. Profiles of structurally submodular separation systems

The most general, or most widely applicable, tree-of-tangles theorem published so far, in the sense of having the weakest premise, is the following:

Theorem 5.4 ([25, Theorem 6])**.** *Let $\vec{S}$ be a submodular separation system in some universe $\vec{U}$ of separations and let $\mathcal{P}$ a set of profiles of S. Then S contains a nested set that distinguishes $\mathcal{P}$.* ⇝ p. 60

The price to pay in Theorem 5.4 for having the mildest set of requirements is that its assertion is also among the weakest of all tree-of-tangles theorems. For graphs, Theorem 5.4 implies only that *for any fixed k* every graph has a tree decomposition displaying its k-tangles. This is a much weaker statement than Theorem 5.1, which finds a tree-decomposition displaying the maximal k-tangles of that graph for all values of k simultaneously.

Let us show how to derive Theorem 5.4 from Lemma 13. For this, let $\mathcal{P}$ be a set of profiles of a submodular separation system S, and for distinct P and P' in $\mathcal{P}$ let

$$\mathcal{A}_{P,P'} := \{\, s \in S : s \text{ distinguishes } P \text{ and } P' \,\}.$$

For proving Theorem 5.4 it suffices to show that the family of all these, i.e. $\mathfrak{A}_{\mathcal{P}} := (\mathcal{A}_{P,P'} : P \neq P' \in \mathcal{P})$, splinters:

Lemma 5.7. *Given a set $\mathcal{P}$ of profiles of a submodular separation system $\vec{S}$, the family $\mathfrak{A}_{\mathcal{P}} = (\mathcal{A}_{P,P'} : P \neq P' \in \mathcal{P})$ splinters.*

Proof. Let $P \neq P'$ and $Q \neq Q'$ be two pairs of profiles in $\mathcal{P}$ and let $r \in \mathcal{A}_{P,P'}$ and $s \in \mathcal{A}_{Q,Q'}$ be two distinct separations. We need to show that we have either $r \in \mathcal{A}_{Q,Q'}$ or $s \in \mathcal{A}_{P,P'}$, or that some corner separation of r and s lies in $\mathcal{A}_{P,P'} \cup \mathcal{A}_{Q,Q'}$. If r and s are nested, then they themselves are corner separations of r and s and there is nothing to show, so let us suppose that r and s cross.

Both r and s are oriented by all four profiles P, P', Q, and Q'. If r distinguishes Q and Q', or if s distinguishes P and P', we are done; so suppose that there are orientations $\vec{r}$ and $\vec{s}$ of r and s with $\vec{r} \in Q \cap Q'$ and $\vec{s} \in P \cap P'$. By possibly switching the roles of P and P', or of Q and Q', we may further assume that $\overleftarrow{r} \in P$ and $\vec{r} \in P'$ as well as $\overleftarrow{s} \in Q$ and $\vec{s} \in Q'$.

The submodularity of S implies that at least one of the two corner separations $\overleftarrow{r} \vee \vec{s}$ and $\vec{r} \vee \overleftarrow{s}$ lies in $\vec{S}$. We will only treat the case that $(\overleftarrow{r} \vee \vec{s}) \in \vec{S}$; the other case is symmetrical.

From the assumption that r and s cross it follows that $\overleftarrow{r} \vee \vec{s}$ is distinct from r and s as an unoriented separation. Therefore, by $\vec{r} \in P'$ and consistency, P' cannot contain $\overleftarrow{r} \vee \vec{s}$ and hence has to contain its inverse $\vec{r} \wedge \overleftarrow{s}$. On the other hand, by $\overleftarrow{r}, \vec{s} \in P$ and the profile property (P), P cannot contain the inverse of $\overleftarrow{r} \vee \vec{s}$ and thus must contain $\overleftarrow{r} \vee \vec{s}$. Now $\overleftarrow{r} \vee \vec{s}$ distinguishes P and P' and is therefore the desired corner separation in $\mathcal{A}_{P,P'}$. □

Let us now deduce Theorem 5.4 from Lemma 13.

Proof of Theorem 5.4. Let $\mathcal{P}$ be a set of profiles of S. By Lemma 5.7 the collection $(\mathcal{A}_{P,P'} : P \neq P' \in \mathcal{P})$ of subsets of S splinters. Each of the $\mathcal{A}_{P,P'}$ is non-empty as P and P' are distinct profiles of S. Thus, by Lemma 13, we can pick one element from each $\mathcal{A}_{P,P'}$ so that the set N of all these elements is a nested set of separations. It is then clear that N distinguishes all the profiles in $\mathcal{P}$. □

The above way of using Lemma 13 to prove a tree-of-tangles theorem is archetypical, and we will use the strategy from this section as a blueprint for the applications of Lemma 13 in the following sections.

5.3.3. Profiles in submodular universes

Theorem 5.4 from the previous section implied that every graph has, for any fixed integer k, a tree-decomposition which displays its k-tangles. However, Robertson's and Seymour's Theorem 5.1 shows that every graph has a tree-decomposition which displays all its *maximal* tangles, i.e., which distinguishes all its distinguishable tangles for all values of k simultaneously, not just for some fixed value of k. Therefore Theorem 5.4 does not imply Theorem 5.1.

Moreover, since Theorem 5.4 does not assume that the universe U it is applied to comes with an order function, Theorem 5.4 cannot say anything about the order of the separations used in the nested set to distinguish all the profiles. If the universe U, as for instance in a graph, *does* come with a submodular order function, one might ask for a nested set which not only distinguishes all the profiles given, but one which does so *efficiently*, i.e., which contains for every pair P, P' of profiles a separation of minimal order among all the separations in U which distinguish P and P'.

The following theorem satisfies both of the requirements above, and is the strongest tree-of-tangles theorem known so far:

Theorem 5.8 (Canonical tree-of-tangles theorem for separation universes [26, Theorem 3.6])**.** *Let $(\vec{U}, \leqslant, {}^*, \vee, \wedge, |\cdot|)$ be a submodular universe of separations. Then* *for every robust set $\mathcal{P}$ of profiles in 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 $T(\mathcal{P}^\alpha) = T(\mathcal{P})^\alpha$; (canonicity)
(iv) if all the profiles in $\mathcal{P}$ are regular, then T is a regular tree set.

Since the definition of robustness of a set of profiles is rather involved, we do not repeat it here. We assure the reader that every set of robust profiles (as defined in Chapter 2) is a robust set of profiles, and that the following proofs robustness will be used only in one place. Therefore we shall use it there as a black box and refer the reader to [26] for the full definition. A very sceptical reader may go forth and read our proofs as if we only made our claims only for sets of robust profiles and not robust sets of profiles.

Since every k-tangle of a graph is robust ([26]), Theorem 5.8 indeed implies Theorem 5.1 of Robertson and Seymour that every graph has a tree-decomposition displaying its maximal tangles (see [26, Section 4.1] for more on building tree-decompositions from nested sets of separations, and how Theorem 5.8 implies Theorem 5.1). Moreover, Theorem 5.8 improves upon Theorem 5.1 by finding a *canonical* such tree-decomposition, i.e., one which is preserved by automorphisms of the graph. Since Lemma 13 does not guarantee any kind of canonicity, we are not able to deduce the full Theorem 5.8 from Lemma 13; however, using Lemma 13 we will be able to find a nested set $T \subseteq U$ with the properties (i), (ii) and (iv). We shall refer to this as the *non-canonical Theorem 5.8.* (In Section 5.4 we shall prove a version of Lemma 13 which implies Theorem 5.8 in full.)

Our strategy will largely be the same as in Section 5.3.2. For a robust set $\mathcal{P}$ of profiles in a submodular universe U we define for every pair P, P' of distinct profiles in $\mathcal{P}$ the set

$$\mathcal{A}_{P,P'} := \{\, a \in U : a \text{ distinguishes } P \text{ and } P' \text{ efficiently} \,\}.$$

Let $\mathfrak{A}_{\mathcal{P}}$ be the family $(\mathcal{A}_{P,P'} : P \neq P' \in \mathcal{P})$. The only lemma we need in order to apply Lemma 13 is the following:

Lemma 5.9. *For a robust set $\mathcal{P}$ of profiles in U the family $\mathfrak{A}_{\mathcal{P}}$ of the sets $\mathcal{A}_{P,P'}$ splinters.*

Proof. Let P, P' and Q, Q' be two pairs of distinguishable profiles in $\mathcal{P}$ and let $r \in \mathcal{A}_{P,P'}$ and $s \in \mathcal{A}_{Q,Q'}$ be two crossing separations. We need to show that we have either $r \in \mathcal{A}_{Q,Q'}$ or $s \in \mathcal{A}_{P,P'}$, or that some corner separation of r and s lies in $\mathcal{A}_{P,P'} \cup \mathcal{A}_{Q,Q'}$. By switching their roles if necessary we may assume that $|r| \leqslant |s|$.

Since Q orients all separations in U of order at most the order of s, Q contains some orientation $\vec{r}$ of r. Similarly Q' contains some orientation of r: if $\overleftarrow{r} \in Q'$ then

r distinguishes Q and Q', and by $|r| \leqslant |s|$ it does so efficiently, giving $r \in \mathcal{A}_{Q,Q'}$. So suppose that $\vec{r} \in Q'$.

If either one of the two corner separations $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \overleftarrow{s}$ has order at most the order of s, then that corner separation would distinguish Q and Q' by the profile property. In particular, that corner separation would do so efficiently and hence lie in $\mathcal{A}_{Q,Q'}$. Thus we may assume that both of these corner separations have order strictly larger than the order of s.

The submodularity of U now implies that both of the other two corner separations, that is, $\vec{r} \wedge \vec{s}$ and $\vec{r} \wedge \overleftarrow{s}$, have order strictly less than the order of r. Therefore both P and P' orient both of these corner separations. By possibly switching the roles of P and P' we may assume that $\overleftarrow{r} \in P$ and $\vec{r} \in P'$. Then P' contains both $\vec{r} \wedge \vec{s}$ and $\vec{r} \wedge \overleftarrow{s}$ due to consistency, since both of these corner separations are distinct from r as unoriented separations by the assumption that r and s cross.

But now the assumption that r distinguishes P and P' efficiently implies that neither of the two corner separations $\vec{r} \wedge \vec{s}$ and $\vec{r} \wedge \overleftarrow{s}$ can distinguish P and P', since the corner separations have strictly lower order than r. Therefore P contains $\vec{r} \wedge \vec{s}$ and $\vec{r} \wedge \overleftarrow{s}$ as well. However, by $\overleftarrow{r} \in P$, this contradicts the robustness of P, which forbids exactly this configuration. □

Let us now deduce the non-canonical Theorem 5.8 from Lemma 13:

Proof of Theorem 5.8, non-canonical. By Lemma 5.9 the collection $\mathfrak{A}_{\mathcal{P}}$ of the sets $\mathcal{A}_{P,P'}$ splinters. Thus by Lemma 13 we can pick an element from each set $\mathcal{A}_{P,P'}$ in $\mathfrak{A}_{\mathcal{P}}$ in such a way that the set T of these elements is nested. Let us show that this set T is as claimed.

For (i), let P and P' be two profiles in $\mathcal{P}$. As T meets the set $\mathcal{A}_{P,P'}$, some element of T distinguishes P and P' by definition of $\mathcal{A}_{P,P'}$.

For (ii), observe that every element of T lies in some $\mathcal{A}_{P,P'}$ and hence distinguishes a pair of profiles in $\mathcal{P}$ efficiently.

Finally, (iv) follows from the fact that all sets $\mathcal{A}_{P,P'}$ in $\mathfrak{A}_{\mathcal{P}}$ are regular if every profiles in $\mathcal{P}$ is regular, which implies that T is a regular tree set in that case. □

5.3.4. Sequences of submodular separation systems

Let us, once more, compare Theorem 5.4 and Theorem 5.8. The first of these has the advantage that it does not depend on any order function and thus applies to a wider class of universes of separations; on the other hand, for those universes that do have an order function, the latter theorem is much more flexible and powerful, since it not only distinguishes all distinguishable profiles across all orders simultaneously, but also does so efficiently.

Our aim in this section is to establish Theorem 14 which combines the advantages of both Theorem 5.4 and Theorem 5.8 (without canonicity), i.e., which is not dependent on the existence of some order function, but which is as powerful and efficient as Theorem 5.8 if such an order function does exist.

Concretely, we shall answer the following question, which inspired this research:

If $S_1 \subseteq S_2 \subseteq \dots \subseteq S_n$ is an ascending sequence of structurally submodular separations systems exhausting a universe of separations U, does there exist a nested set of separations which efficiently distinguishes all the maximal profiles in U?

Let us substantiate this question with rigorous definitions of the terms involved.

We call a sequence $S_1 \subseteq S_2 \subseteq \dots \subseteq S_n \subseteq U$ of submodular separation systems in a universe U *compatible* if for all pairs $s_i \in S_i$ and $s_j \in S_j$ with $i \leqslant j$, either S_i contains at least two corner separations of s_i and s_j, or S_j contains at least three corner separations of s_i and s_j.

Observe that if U comes with a submodular order function $|\cdot|$ and the S_i are defined as in Section 5.3.3, i.e., if S_i is the set of all separations in U of order $< i$, then the sequence $S_1 \subseteq S_2 \subseteq \dots \subseteq S_n \subseteq U$ is a compatible sequence of submodular separation systems.

A *profile in* $\mathcal{S} = (S_1, \dots, S_n)$ is a profile of one of the S_i.

A separation $s \in S_n$ *distinguishes* two profiles P and Q in $\mathcal{S}$ if there are orientations of s such that $\vec{s} \in P$ and $\overleftarrow{s} \in Q$. The separation s distinguishes P and Q *efficiently* if $s \in S_i$ for every S_i which contains a separation that distinguishes P and Q.

Note once more that, as above, these notions of profiles and efficient distinguishers coincide with their usual definitions as given in Section 5.3.3 if U has a submodular order function and the S_i are the subsets of U containing all separations of order $< i$.

We also require a structural formulation of the concept of robustness from [26]: A set $\mathcal{P}$ of profiles in $\mathcal{S}$ is *robust* if for all $P, Q, Q' \in \mathcal{P}$ the following holds: for every $\vec{r} \in Q \cap Q'$ with $\overleftarrow{r} \in P$ and every s which distinguishes Q and Q' efficiently, if $s \in S_j$, then there is an orientation $\vec{s}$ of s such that either $(\overleftarrow{r} \vee \vec{s}) \in P$ or $(\vec{r} \vee \vec{s}) \in \vec{S_j}$.

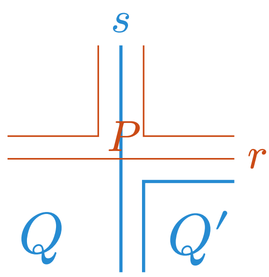

Figure 5.1.: Robustness.

With the above definitions we are now able to formally state and prove Theorem 14, which includes both Theorem 5.4 and the non-canonical Theorem 5.8 (and hence Theorem 5.1) as special cases:

Theorem 14. *If $\mathcal{S} = (S_1, \dots, S_n)$ is a compatible sequence of structurally submodular separation systems inside a universe U, and $\mathcal{P}$ is a robust set of profiles in $\mathcal{S}$, then there is a nested set N of separations in U which efficiently distinguishes all the distinguishable profiles in $\mathcal{P}$.*

Since the proof of Theorem 14 runs along very similar lines as the proof of Theorem 5.8 in the previous section we only sktech it here:

Sketch of proof. For every pair P, P' of distinguishable profiles in $\mathcal{P}$ let $\mathcal{A}_{P,P'}$ be the set of all $s \in S_n$ that distinguish P and P' efficiently. The assertion of Theorem 14 follows directly from Lemma 13 if we can show that the family $\mathfrak{A}$ of these sets $\mathcal{A}_{P,P'}$ splinters.

So let $r \in \mathcal{A}_{P,P'}$ and $s \in \mathcal{A}_{Q,Q'}$ be given. If r and s are nested there is nothing to show, so suppose they cross. Let i and j be minimal integers such that $r \in S_i$ and $s \in S_j$; we may assume without loss of generality that $i \leqslant j$.

If r distinguishes Q and Q', then $r \in \mathcal{A}_{Q,Q'}$, so suppose not, that is, suppose that some orientation $\vec{r}$ of r lies in both Q and Q'.

If one of the two corner separations $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \overleftarrow{s}$ lies in $\vec{S_j}$, then that separation distinguishes Q and Q' by consistency and the profile property and hence would lie in $\mathcal{A}_{Q,Q'}$. So we may suppose that neither of these two corner separations lies in $\vec{S_j}$. The compatibility of $\mathcal{S}$ then implies that both of the other two corner separations, $\overleftarrow{r} \vee \vec{s}$ and $\overleftarrow{r} \vee \overleftarrow{s}$, lie in S_i.

By possibly switching the roles of P and P' we may assume that $\vec{r} \in P'$ and $\overleftarrow{r} \in P$. Then the robustness of $\mathcal{P}$ implies that P contains either $\overleftarrow{r} \vee \vec{s}$ or $\overleftarrow{r} \vee \overleftarrow{s}$. This corner separation then lies in $\mathcal{A}_{P,P'}$ due to the consistency of P'. $\square$

Theorem 14 directly implies both Theorem 5.4 and the non-canonical Theorem 5.8: for the first theorem, consider the singleton sequence $S_1 = S$; and for the latter, take as S_i the set of all separations of order $< i$ and let n be large enough that $S_n = U$.

5.3.5. Clique-separations in finite graphs

For a finite graph G a separation (A, B) of G is a *clique separation* if the induced subgraph $G[A \cap B]$ is a complete graph. Clique separations in graphs have been studied by various people over the course of the last century [53, 68]. More recently clique separations have received quite some attention in theoretical computer science (see for instance [2, 15, 60]) following Tarjan's work [66] on their algorithmic aspects.

In [25] it was shown that the theory of submodular separation systems can be applied to clique separations of finite graphs to deduce the existence of certain nested distinguishing sets. Using Lemma 13 directly instead of Theorem 5.4, we are able to obtain a stronger result than the one given in [25], much in the same way that Theorem 14 improves upon Theorem 5.4.

Given a finite graph $G = (V, E)$, let $\vec{U} = \vec{S}(G)$ be the universe of separations of G and let $\vec{\mathcal{K}} = \vec{\mathcal{K}}(G) \subseteq \vec{U}$ be the separation system of all clique separations of G. Consequently, $\vec{\mathcal{K}}_k = \vec{\mathcal{K}}_k(G)$ is the set of all clique-separations in G of order less than k, i.e., the set of all $(A, B) \in \vec{\mathcal{K}}$ such that $|A \cap B| < k$.

It was shown in [25, Lemma 17] that this $\mathcal{K}$ is a submodular separation system. Following their proof, we can show that in fact every such $\mathcal{K}_k \subseteq \mathcal{K}$ is submodular in U, and that these extend each other in a way similar to the ordinary S_k of vertex separations of G:

Lemma 5.10. *Let r and s be two crossing clique separations with $|r| \leqslant |s|$. Then there are orientations $\vec{r}$ and $\vec{s}$ of r and s such that $(\overleftarrow{r} \wedge \overleftarrow{s})$, $(\overleftarrow{r} \wedge \vec{s})$, and $(\vec{r} \wedge \overleftarrow{s})$ are clique separations with $|\overleftarrow{r} \wedge \overleftarrow{s}| \leqslant |r|$ and $|\overleftarrow{r} \wedge \vec{s}| \leqslant |r|$ as well as $|\vec{r} \wedge \overleftarrow{s}| \leqslant |s|$. Moreover, if $|\overleftarrow{r} \wedge \overleftarrow{s}| = |r| = |s|$, then $(\vec{r} \wedge \vec{s})$ is also a clique separation with $|\vec{r} \wedge \vec{s}| \leqslant |r|$.*

Proof. Let $s = \{ A, B \}$ and $t = \{ C, D \}$ be two crossing clique separations of G with $|r| \leqslant |s|$. Since $C \cap D$ is a separator of G, and all vertices in $A \cap B$ are pairwise adjacent, $A \cap B$ must be a subset of either C or D. Similarly $C \cap D$ must be a subset of either A or B. By renaming the sets if necessary we may assume that $A \cap B \subseteq C$ and $C \cap D \subseteq A$. We orient r as $\vec{r} = (A, B)$ and s as $\vec{s} = (C, D)$; let us show that these orientations are as claimed.

Observe first that the separators of both $(\overleftarrow{r} \wedge \overleftarrow{s})$ and $(\overleftarrow{r} \wedge \vec{s})$ are subsets of $A \cap B$, showing that these are clique separations of order at most $|r| = |A \cap B|$. Similarly, the separator of the corner separation $(\vec{r} \wedge \overleftarrow{s})$ is a subset of $C \cap D$, and hence $(\vec{r} \wedge \overleftarrow{s})$ is a clique separation of order at most $|s| = |C \cap D|$.

Finally, suppose that $|\overleftarrow{r} \wedge \overleftarrow{s}| = |r| = |s|$. Then, since the separator of $(\overleftarrow{r} \wedge \overleftarrow{s})$ is a subset of both $A \cap B$ and of $C \cap D$, this separator must in fact be equal to both $A \cap B$ and $C \cap D$. Consequently the separator of $(\vec{r} \wedge \vec{s})$ also equals $A \cap B = C \cap D$, which shows that $(\vec{r} \wedge \vec{s})$ is a clique separation of order at most r. $\square$

We can now consider *clique-profiles in G* with respect to these separation system, that is a profile P of order k is a consistent orientation of $\mathcal{K}_k$ satisfying the profile property

$$\forall\, \vec{r}, \vec{s} \in P \colon (\overleftarrow{r} \wedge \overleftarrow{s}) \notin P\,. \tag{P}$$

Every hole in G (i.e., an induced cycle of length at least 4) defines a clique-profile P of order $|V|$ in G by letting P contain a separation $(A, B) \in \vec{\mathcal{K}}$ of order less than $|V|$ if and only if that hole is contained in $G[B]$. In an analogous way every clique of size k defines a clique-profile of order k in G. Let us denote by $\mathcal{P}_k$ the set of all clique-profiles of order k.

As usual, given two distinguishable profiles P and P', let

$$\mathcal{A}_{P,P'} := \{\, a \in \mathcal{K} : a \text{ distinguishes } P, P' \text{ efficiently} \,\}.$$

We will show that the collection of these $\mathcal{A}_{P,P'}$ splinters.

Lemma 5.11. *For any set $\mathcal{P}$ of clique-profiles the collection*

$$(\mathcal{A}_{P,P'} : P, P' \text{ distinguishable profiles in } \mathcal{P})$$

splinters.

Proof. Let P, P' and Q, Q' be two pairs of distinguishable profiles in $\mathcal{P}$ and let $r \in \mathcal{A}_{P,P'}$ and $s \in \mathcal{A}_{Q,Q'}$ be two distinct separations. We need to show that we have either $r \in \mathcal{A}_{Q,Q'}$ or $s \in \mathcal{A}_{P,P'}$, or that some corner separation of r and s lies in $\mathcal{A}_{P,P'} \cup \mathcal{A}_{Q,Q'}$. If r and s are nested, then the latter is immediate, so suppose that r and s cross. By switching their roles if necessary we may further assume that $|r| \leqslant |s|$.

Since Q orients s, and $|r| \leqslant |s|$, the profile Q contains some orientation $\vec{r}$ of r. Similarly Q' contains some orientation of r. If $\overleftarrow{r} \in Q'$, then r distinguishes Q and Q', and by $|r| \leqslant |s|$ it does so efficiently, giving $r \in \mathcal{A}_{Q,Q'}$. So suppose that $\vec{r} \in Q'$.

By Lemma 5.10 at least three of the corner separations of r and s are clique separations of order at most $|s|$. Thus at least one of $(\overleftarrow{r} \wedge \overleftarrow{s})$ and $(\overleftarrow{r} \wedge \vec{s})$ is a clique separation of order at most $|s|$. This corner separation then distinguishes Q and Q' by the profile property, and in fact it does so efficiently, since its order is at most $|s|$, yielding the desired corner separation in $\mathcal{A}_{Q,Q'}$. □

It is now straightforward to use Lemma 13 to obtain the following theorem:

Theorem 15. *There is a nested set of separations which efficiently distinguishes all the distinguishable clique-profiles in $\bigcup_{i=1}^{n} \mathcal{P}_i$.*

Proof. By Lemma 5.11, we can apply Lemma 13 to

$$(\mathcal{A}_{P,P'} : P, P' \text{ distinguishable profiles in } \bigcup_i \mathcal{P}_i),$$

resulting in the claimed nested subset. □

In particular, for any two holes, a hole and a clique, or two cliques if there is a clique separation which distinguishes them, then our nested set contains one such separation of minimal order. As usual, such a nested set can be transformed into a tree-decomposition of G (see [21] for details). Thus G admits a tree-decomposition whose adhesion sets are cliques and which efficiently distinguishes all the holes and cliques distinguishable by clique separations in G. Such a decomposition is similar to, but not exactly the same as, the decomposition constructed by R. E. Tarjan [66].

We will see in Section 5.5.2 that such a decomposition can in fact be chosen canonically, i.e., to invariantly under automorphisms of G.

5.4. The Canonical Splinter Lemma

As we saw in the previous section, Lemma 13 is already strong enough to imply most of Theorem 5.8, but crucially does not guarantee the canonicity asserted in (iii). In this section we wish to prove a version of Lemma 13 using a stronger set of assumptions from which we can deduce Theorem 5.8 in full: we want to find, for a family $\mathfrak{A} = (\mathcal{A}_i : i \in I)$ of subsets of some universe U, a nested set $N = N(\mathfrak{A})$ meeting all the $\mathcal{A}_i$ that is *canonical*, i.e., which only depends on invariants of $\mathfrak{A}$. More formally, we want to find $N = N(\mathfrak{A})$ in such a way that if $\mathfrak{A}' = (\mathcal{A}'_i : i \in I)$ is another family of subsets of some other universe U' that also meets the assumptions of our theorem, and φ is an isomorphism of separation systems between $\bigcup_{i\in I} \mathcal{A}_i$ and $\bigcup_{i\in I} \mathcal{A}'_i$ with $\varphi(\mathcal{A}_i) = \mathcal{A}'_i$ for all $i \in I$, we ask that $N(\mathfrak{A}') = \varphi(N(\mathfrak{A}))$. In particular, the nested set found by our theorem should not depend on the universe into which the family $\mathfrak{A}$ is embedded.

The assumptions of Lemma 13 are not sufficient to guarantee the existence of such a canonical set. Consider the example where we have just two separations, s and t, which are crossing and let $\mathfrak{A} = (\mathcal{A}_1) = (\{ s, t \})$. Note that $\mathfrak{A}$ splinters, but there may be an automorphism that swaps the two separations so the choice of any single one of them is non-canonical. Since the separations are crossing we cannot use both of them for our nested set either.

For obtaining a canonical nested set, one crucial ingredient will be the notion of *extremal* elements of a set of separations, which was already used in [26]. Given a set $A \subseteq U$ of (unoriented) separations, an element $a \in A$ is *extremal* in A, or an *extremal element* of A, if a has some orientation $\vec{a}$ that is a maximal element of $\vec{A}$. (Recall that $\vec{A}$ is the set of orientations of separations in A.) The set of extremal elements of a set of separations is an invariant of separation systems in the following sense: if E is the set of extremal elements of some set $A \subseteq S$ of separations, and φ is an isomorphism between $\vec{S}$ and some other separation system, then $\varphi(E)$ is precisely the set of extremal separations of $\varphi(A)$. Moreover, the extremal separations of a set $A \subseteq U$ are nested with each other under relatively weak assumptions: for instance, it suffices that for any two separations in A at least two of their corner separations also lie in A.

Let us formally state a set of assumptions under which we can prove a canonical version of Lemma 13. Given two separations r and s and two of their corner separations c_1 and c_2, we say that c_1 and c_2 are *from different sides of* r if, for orientations of c_1, r, and s with $\vec{c}_1 = (\vec{r} \wedge \vec{s})$, there is an orientation $\vec{c}_2$ of c_2 such that either $\vec{c}_2 = (\overleftarrow{r} \wedge \vec{s})$ or $\vec{c}_2 = (\overleftarrow{r} \wedge \overleftarrow{s})$. Note that c_1 and c_2 being from different sides of r does not imply that c_1 and c_2 are distinct separations; consider for instance the edge case that $r = s = c_1 = c_2$.

Let $\mathfrak{A} = (\mathcal{A}_i : i \in I)$ be a finite collection of non-empty finite subsets of U and let $\preccurlyeq$ be any partial order on I. We write $i \prec j$ if and only if $i \preccurlyeq j$ and $i \neq j$. We say that $\mathfrak{A}$ *splinters hierarchically* if for all $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$ the following two conditions hold:

(HS1) If $i \prec j$, either some corner separation of a_i and a_j lies in $\mathcal{A}_j$, or two corner separations of a_i and a_j from different sides of a_i lie in $\mathcal{A}_i$.

(HS2) If neither $i \prec j$ nor $j \prec i$, there are $k \in \{ i, j \}$ and corner separations c_1 and c_2 of a_i and a_j from different sides of a_k such that $c_1 \in \mathcal{A}_k$ and $c_2 \in \mathcal{A}_i \cup \mathcal{A}_j$.

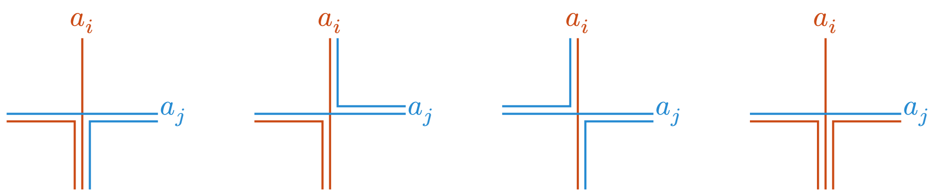

Figure 5.2.: The possible configurations in condition (HS2), up to symmetry.

In particular if $\preccurlyeq$ is the trivial partial order on I in which all $i \neq j$ are incomparable, then $\mathfrak{A}$ splinters hierarchically if and only if (HS2) holds for all $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$; this special case which ignores the partial order on I is perhaps the cleanest form of an assumption that suffices for a canonical nested set meeting all $\mathcal{A}_i$ in $\mathfrak{A}$. The reason we need to allow a partial order $\preccurlyeq$ on I and the slightly weaker condition in (HS1) for comparable elements of I is that otherwise we would not be able to deduce Theorem 5.8 in full from our main theorem of this section due to a quirk in the way that robustness is defined for profiles in [26] (see Section 5.5).

Our first lemma enables us to find a canonical nested set inside $\bigcup_{i\in I} \mathcal{A}_i$ for a collection of sets $\mathcal{A}_i$ whose indexing set is an antichain:

Lemma 5.12. *Let $(\mathcal{A}_i : i \in I)$ be a collection of subsets of U that splinters hierarchically. If $K \subseteq I$ is an antichain in $\preccurlyeq$, then the set of extremal elements of $\bigcup_{k\in K} \mathcal{A}_k$ is nested.*

Proof. Suppose that $K \subseteq I$ is an antichain and that for some $i, j \in K$ there are $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$ such that a_i and a_j are extremal in $\bigcup_{k\in K} \mathcal{A}_k$ but cross. Let $\vec{a}_i$ and $\vec{a}_j$ be the orientations of a_i and a_j witnessing their extremality. As a_i and a_j cross, there are three ways of orienting a_i and a_j so that the supremum of this orientation is strictly larger than $\vec{a}_i$ or $\vec{a}_j$. Hence none of these corner separations can lie in $\mathcal{A}_i \cup \mathcal{A}_j$, since that would contradict the maximality of $\vec{a}_i$ or $\vec{a}_j$ in $\bigcup_{k\in K} \vec{\mathcal{A}}_k$. On the other hand, since neither $i \prec j$ nor $j \prec i$, by condition (HS2) and the assumption that a_i and a_j cross there are at least two orientations of a_i and a_j whose corresponding supremum lies in $\mathcal{A}_i \cup \mathcal{A}_j$, causing a contradiction to the extremality of a_i and a_j. $\square$

We are now able to prove a canonical version of the Splinter Lemma by repeatedly applying Lemma 5.12 to the collection of the $\mathcal{A}_i$ of $\preccurlyeq$-minimal index that have not yet been met by the nested set constructed so far:

Lemma 16 (Canonical splinter lemma)**.** *Let U be a universe of separations and let $\mathfrak{A} = (\mathcal{A}_i : i \in I)$ be a collection of subsets of U that splinters hierarchically with respect to a partial order $\preccurlyeq$ on I. Then there exists a nested set $N = N(\mathfrak{A})$ meeting every $\mathcal{A}_i$ in $\mathfrak{A}$.*

Moreover, $N(\mathfrak{A})$ is canonical: if φ is an isomorphism of separation systems between $\bigcup_{i\in I} \vec{\mathcal{A}}_i$ and a subset of some universe U' such that the family $\varphi(\mathfrak{A}) := (\varphi(\mathcal{A}_i) : i \in I)$ splinters hierarchically with respect to $\preccurlyeq$, then $N(\varphi(\mathfrak{A})) = \varphi(N(\mathfrak{A}))$.

Proof. We proceed by induction on $|I|$. If $|I| = 1$ we can choose as N the set of extremal elements of $\mathcal{A}_i$, which is nested by Lemma 5.12 and clearly canonical.

So suppose that $|I| > 1$ and that the claim holds for all smaller index sets. Let K be the set of minimal elements of I with respect to $\preccurlyeq$. By Lemma 5.12 the set $E = E(\mathfrak{A})$ of extremal elements of $\bigcup_{k\in K} \mathcal{A}_k$ is nested. Let $J \subseteq I$ be the set of indices of all those $\mathcal{A}_j$ that do not meet E, and for $j \in J$ let $\mathcal{A}'_j$ be the set of all elements of $\mathcal{A}_j$ that are nested with E. We claim that the collection $\mathfrak{A}' = (\mathcal{A}'_j : j \in J)$ splinters hierarchically with respect to $\preccurlyeq$ on J. This follows from Lemma 2.1 as soon as we show that each $\mathcal{A}'_j$ is non-empty.

To see that each $\mathcal{A}'_j$ is non-empty, for $j \in J$ let a_j be an element of $\mathcal{A}_j$ that crosses as few elements of E as possible. We wish to show that a_j is nested with E and thus $a_j \in \mathcal{A}'_j$. So suppose that a_j crosses some separation in E, that is, some $a_i \in \mathcal{A}_i \cap E$ with $i \in I \setminus J$. Since i is a minimal element of I we have either $i \preccurlyeq j$ or that i and j are incomparable. We shall treat these cases separately.

Consider first the case that $i \prec j$. By condition (HS1) of splintering hierarchically, either some corner separation of a_i and a_j lies in $\mathcal{A}_j$, or two corner separations of a_i and a_j from different sides of a_i lie in $\mathcal{A}_i$. The first of these possibilities contradicts the choice of a_j, since that corner separation in $\mathcal{A}_j$ would cross fewer elements of E by Lemma 2.1. On the other hand, the latter of these possibilities contradicts the choice of a_i as an extremal element of $\bigcup_{k\in K} \mathcal{A}_k$. Thus the case $i \preccurlyeq j$ is impossible.

Let us now consider the case that i and j are incomparable. Again, by the choice of a_j, none of the corner separations of a_i and a_j can lie in $\mathcal{A}_j$ by Lemma 2.1. Therefore condition (HS2) of splintering hierarchically yields the existence of a corner separation of a_i and a_j in $\mathcal{A}_i$ for each side of a_i; this, however, contradicts the extremality of a_i in $\bigcup_{k\in K} \mathcal{A}_k$ as before.

Therefore each of the sets $\mathcal{A}'_j$ with $j \in J$ is non-empty, and hence the collection $\mathfrak{A}' = (\mathcal{A}'_j : j \in J)$ splinters hierarchically with respect to $\preccurlyeq$. Since $|J| < |I|$ we may apply the induction hypothesis to this collection to obtain a canonical nested set $N' = N(\mathfrak{A}')$ meeting all $\mathcal{A}'_j$. Now $N = N' \cup E$ is a nested subset of U which meets every $\mathcal{A}_i$ for $i \in I$. It remains to show that N is canonical.

To see that N is canonical let φ be an isomorphism of separation systems between $\bigcup_{i\in I}\vec{\mathcal{A}}_i$ and a subset of some universe U' such that $\varphi(\mathfrak{A})$ splinters hierarchically with respect to $\preccurlyeq$ in U'. Then $\varphi(E) = E(\varphi(\mathfrak{A}))$, i.e., the set of extremal elements of $\bigcup_{i\in I}\varphi(\mathcal{A}_i)$ is exactly $\varphi(E)$. Therefore $\varphi(E)$ meets $\varphi(\mathcal{A}_i)$ if and only if E meets $\mathcal{A}_i$. Consequently the restriction of φ to $\bigcup_{j\in J}\vec{\mathcal{A}}_j{}'$ is an isomorphism of separation systems between $\bigcup_{j\in J}\vec{\mathcal{A}}_j{}'$ and its image in U' with the property that $\varphi(\mathfrak{A}')$ splinters hierarchically with respect to $\preccurlyeq$ on J. Moreover, for $j \in J$, the image $\varphi(\mathcal{A}_j')$ of $\mathcal{A}_j'$ is exactly the set of those separations in $\varphi(\mathcal{A}_j)$ that are nested with $\varphi(E)$.

Thus we can apply the induction hypothesis to find that $N(\varphi(\mathfrak{A}')) = \varphi(N(\mathfrak{A}'))$. Together with the above observation that $\varphi(E(\mathfrak{A})) = E(\varphi(\mathfrak{A}))$ this gives

$$\varphi(N(\mathfrak{A})) = \varphi(E(\mathfrak{A})) \cup \varphi(N(\mathfrak{A}')) = E(\varphi(\mathfrak{A})) \cup N(\varphi(\mathfrak{A}')) = N(\varphi(\mathfrak{A})),$$

concluding the proof. $\square$

5.5. Applications of the Canonical Splinter Lemma

In this section we apply Lemma 16 to obtain a short proof of Theorem 5.8, to strengthen Theorem 15 for clique separations so as to make it canonical, and finally to establish a canonical tree-of-tangles theorem for another type of separations, so-called circle separations.

5.5.1. Robust profiles

Having established Lemma 16 in the previous section, we are now ready to derive the full version of Theorem 5.8. For this let $(\vec{U}, \leqslant, {}^*, \vee, \wedge, |\cdot|)$ be a submodular universe of separations and $\mathcal{P}$ a robust set of profiles in U, and let I be the set of all pairs of distinguishable profiles in $\mathcal{P}$. As in Section 5.3.3, for $\{P, P'\} \in I$ we let

$$\mathcal{A}_{P,P'} := \{\, a \in U : a \text{ distinguishes } P \text{ and } P' \text{ efficiently} \,\},$$

and let $\mathfrak{A}_{\mathcal{P}}$ be the family $(\mathcal{A}_{P,P'} : \{P, P'\} \in I)$. We furthermore define a partial order $\preccurlyeq$ on I by letting $\{P, P'\} \prec \{Q, Q'\}$ if and only if the order of some element of $\mathcal{A}_{P,P'}$ is strictly lower than the order of some element of $\mathcal{A}_{Q,Q'}$. Note that the separations in a fixed $\mathcal{A}_{P,P'}$ all have the same order.

We shall be able to deduce Theorem 5.8 from Lemma 16 as soon as we show that $\mathfrak{A}_{\mathcal{P}}$ splinters hierarchically.

Lemma 5.13. *$\mathfrak{A}_{\mathcal{P}}$ splinters hierarchically with respect to $\preccurlyeq$.*

Proof. Let $r \in \mathcal{A}_{P,P'}$ and $s \in \mathcal{A}_{Q,Q'}$ be given. By switching their roles if necessary we may assume that $|r| \leqslant |s|$. Then Q and Q' both orient r; we may assume without loss of generality that $\vec{r} \in Q$. We will make a case distinction depending on the way Q' orients r.

Let us first treat the case that Q and Q' orient r differently, i.e., that $\overleftarrow{r} \in Q'$. Then r distinguishes Q and Q' and hence $|r| = |s|$ by the efficiency of s. This implies that $\{ P, P' \}$ and $\{ Q, Q' \}$ are either the same pair or else incomparable in $\preccurlyeq$. We may assume further without loss of generality that $\vec{s} \in Q$ and $\overleftarrow{s} \in Q'$. Consider now the two corner separations $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$: if at least one of these two has order at most $|s|$, then this corner separation would distinguish Q and Q' by the profile property. The efficiency of s would then imply that this corner separation has order exactly $|s|$ and hence lies in $\mathcal{A}_{Q,Q'}$. The submodularity of the order function implies that this is the case for at least one, and therefore for both of these corner separations, yielding the existence of two corner separations of r and s from different sides of s in $\mathcal{A}_{Q,Q'}$ and showing that (HS2) is satisfied.

Let us now consider the case that Q and Q' orient r in the same way, i.e., that $\vec{r} \in Q'$. We make a further split depending on whether $|r| = |s|$ or $|r| < |s|$.

Suppose first that $|r| = |s|$; then neither $\{ P, P' \} \prec \{ Q, Q' \}$ nor $\{ Q, Q' \} \prec \{ P, P' \}$. We may assume that P and P' orient s in the same way: for if P and P' orient s differently, we may switch the roles of r and s as well as $\{ P, P' \}$ and $\{ Q, Q' \}$ and apply the above case. So suppose that both of P and P' contain $\vec{s}$, say. Then neither of the corner separations $\vec{r} \vee \overleftarrow{s}$ nor $\overleftarrow{r} \vee \vec{s}$ can have order strictly less than $|r| = |s|$, as these corner separations would distinguish Q and Q' or P and P', respectively, and would therefore contradict the efficiency of s or of r, respectively. The submodularity of $|\cdot|$ now implies that both of these corner separations have order exactly $|r| = |s|$ and hence lie in $\mathcal{A}_{Q,Q'}$ and $\mathcal{A}_{P,P'}$, respectively, showing that (HS2) holds.

Finally, let us suppose that $|r| < |s|$; then $\{ P, P' \} \prec \{ Q, Q' \}$. Consider the two corner separations $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \vec{s}$: if both of $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \vec{s}$ have order strictly greater than $|s|$, then by the submodularity of the order function both of the other two corner separations $\overleftarrow{r} \vee \vec{s}$ and $\overleftarrow{r} \vee \overleftarrow{s}$ have order strictly smaller than $|r|$. By the robustness of $\mathcal{P}$ one of these two corner separations would distinguish P and P', contradicting the efficiency of r.

Thus we may assume at least one of $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \vec{s}$ has order at most $|s|$. Then that corner separation distinguishes Q and Q'. In fact, it does so efficiently and hence lies in $\mathcal{A}_{Q,Q'}$, showing that (HS1) holds and concluding the proof. □

We are now ready to deduce the full Theorem 5.8 from Lemma 16:

Theorem 5.8 (Canonical tree-of-tangles theorem for separation universes [26, Theorem 3.6]). *Let $(\vec{U}, \leqslant, {}^*, \vee, \wedge, |\cdot|)$ be a submodular universe of separations. Then for every robust set $\mathcal{P}$ of profiles in U there is a nested set $T = T(\mathcal{P}) \subseteq U$ of separations such that:* ⇝ p. 68

(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 $T(\mathcal{P}^\alpha) = T(\mathcal{P})^\alpha$; (canonicity)

(iv) if all the profiles in $\mathcal{P}$ are regular, then T is a regular tree set.

Proof. By Lemma 5.13 the family $\mathfrak{A}_{\mathcal{P}}$ splinters hierarchically. Thus we can apply Lemma 16 to $\mathfrak{A}_{\mathcal{P}}$ to obtain a nested set $N = N(\mathfrak{A}_{\mathcal{P}})$ which meets every $\mathcal{A}_{P,P'}$. Clearly, N satisfies (i), (ii) and (iv) of Theorem 5.8.

To see that N satisfies (iii), let α be an automorphism of $\vec{U}$. Then the restriction of α to $\bigcup_{\{P,P'\}\in I} \vec{\mathcal{A}}_{P,P'}$ is an isomorphism of separation systems onto its image in $\vec{U}$. We therefore have, by Lemma 16, that $\alpha(N(\mathfrak{A}_{\mathcal{P}})) = N(\alpha(\mathfrak{A}_{\mathcal{P}}))$. For every $\mathcal{A}_{P,P'}$ in $\mathfrak{A}_{\mathcal{P}}$ we have that $\alpha(\mathcal{A}_{P,P'})$ is precisely the set of those separations in U which distinguish P^α and P'^α efficiently; in other words, we have $\alpha(\mathfrak{A}_{\mathcal{P}}) = \mathfrak{A}_{\mathcal{P}^\alpha}$, showing that (iii) is satisfied. $\square$

5.5.2. Clique separations

Regarding the clique-profiles discussed in Section 5.3.5, Lemma 5.10 not only suffices to show that the sets $\mathcal{A}_{P,P'}$ splinter, but can be used to show that the collection of these $\mathcal{A}_{P,P'}$ even splinters hierarchically, allowing us to apply Lemma 16: for this we simply define the same partial order $\preccurlyeq$ on the set of pairs $\{P, P'\}$ as in the previous section, that is, $\{P, P'\} \prec \{Q, Q'\}$ if and only if $|r| < |s|$ for some (equivalently: for all) $r \in \mathcal{A}_{P,P'}$ and $s \in \mathcal{A}_{Q,Q'}$.

To see this, let P, P' and Q, Q' be distinguishable pairs of profiles of clique separations. Let $r \in \mathcal{A}_{P,P'}$ and $s \in \mathcal{A}_{Q,Q'}$, and suppose without loss of generality that $|r| \leqslant |s|$. If r and s are nested, then r and s themselves are corner separations of r and s that lie in $\mathcal{A}_{P,P'}$ and $\mathcal{A}_{Q,Q'}$, respectively. However, if r and s cross, then by Lemma 5.10 there are orientations of r and s such that $|\overleftarrow{r} \wedge \overleftarrow{s}|, |\overleftarrow{r} \wedge \vec{s}| \leqslant |r|$ and $|\overleftarrow{r} \wedge \overleftarrow{s}|, |\overleftarrow{r} \wedge \vec{s}|, |\vec{r} \wedge \overleftarrow{s}| \leqslant |s|$. By switching their roles if necessary we may assume that $\vec{r} \in P$ and $\overleftarrow{r} \in P'$, and likewise that $\vec{s} \in Q$ and $\overleftarrow{s} \in Q'$.

Since $(\vec{r} \wedge \overleftarrow{s}), (\overleftarrow{r} \wedge \overleftarrow{s}) \leqslant \overleftarrow{s}$ and $\overleftarrow{s} \in Q'$, the profile Q' contains both of these corner separations by consistency. On the other hand, by the assumption that $|r| \leqslant |s|$, the separation r gets oriented by Q, and consequently, by the profile property, Q must contain the inverse of one of those two corner separations. This corner separation then distinguishes Q and Q', and in fact it does so efficiently, since its order is at most $|s|$, meaning that this corner separation lies in $\mathcal{A}_{Q,Q'}$. Therefore, if $|r| < |s|$, then condition (HS1) of splintering hierarchically is satisfied.

So suppose further that $|r| = |s|$, and let us check that condition (HS2) of splintering hierarchically is satisfied. Observe that, similarly as above, P orients s, and P' contains both $(\overleftarrow{r} \wedge \overleftarrow{s})$ and $(\overleftarrow{r} \wedge \vec{s})$ by consistency with $\overleftarrow{r} \in P'$, implying as before that one of $(\overleftarrow{r} \wedge \overleftarrow{s})$ and $(\overleftarrow{r} \wedge \vec{s})$ also efficiently distinguishes P and P', i.e., is an element of $\mathcal{A}_{P,P'}$. If this corner separation in $\mathcal{A}_{P,P'}$ and the corner separation in $\mathcal{A}_{Q,Q'}$ found above are from different sides of either r or s, then condition (HS2) of splintering hierarchically would be satisfied. So suppose not; that is, suppose that $(\overleftarrow{r} \wedge \overleftarrow{s})$ distinguishes both P and P' as well as Q and Q' efficiently.

In particular $|\overleftarrow{r} \wedge \overleftarrow{s}| = |r| = |s|$, and hence by the last part of Lemma 5.10, all four corner separations of r and s have order at most $|r|$. Consequently, since P' orients s, one of $(\vec{r} \wedge \overleftarrow{s})$ and $(\vec{r} \wedge \vec{s})$ distinguishes P and P' efficiently, which one depending on whether $\vec{s} \in P'$ or $\overleftarrow{s} \in P'$. In either case we have found a corner separation of r and s in $\mathcal{A}_{P,P'}$, which together with $(\overleftarrow{r} \wedge \overleftarrow{s}) \in \mathcal{A}_{Q,Q'}$ witnesses that (HS2) is fulfilled.

Therefore, by Lemma 16 we get that we can choose the set in Theorem 15 canonically:

Theorem 17. *For every set $\mathcal{P}$ of clique-profiles of a graph G, there is a nested set $N = N(\mathcal{P})$ of separations which efficiently distinguishes all the distinguishable clique-profiles in $\mathcal{P}$ and is canonical, that is, such that $N(\mathcal{P}^\alpha) = N(\mathcal{P})^\alpha$ for every automorphism α of the underlying graph G.*

Proof. Every automorphism of G induces an automorphism of the separation system. Hence we can obtain the claimed nested set by applying Lemma 16 to the family of the sets $\mathcal{A}_{P,P'}$ of those clique separations which efficiently distinguish the pair P, P' of distinguishable profiles in $\mathcal{P}$. $\square$

5.5.3. Circle separations

Another special case of separation systems are those of *circle separations* discussed in [25]: given a fixed cyclic order on a ground-set V, a *circle separation* of V is a bipartition (A, B) of V into two disjoint intervals in the cyclic order. Observe that the set of all circle separations is not closed under joins and meets and hence not a sub-universe of the universe of all bipartitions of V:

Example 5.14. Consider the natural cyclic order on the set $V = \{1, 2, 3, 4\}$. The bipartitions $(\{1\}, \{2, 3, 4\})$ and $(\{3\}, \{4, 1, 2\})$ of V are circle separations. However, their supremum in the universe of all bipartitions of V is $(\{1, 3\}, \{2, 4\})$, which is not a circle separation.

Let V be a ground-set with a fixed cyclic order and $(\vec{U}, \leqslant, {}^*, \vee, \wedge, |\cdot|)$ the universe of all bipartitions of V with any submodular order function $|\cdot|$. Let $\vec{S} \subseteq \vec{U}$ be the set of all separations in $\vec{U}$ that are circle separations of V. We denote by $\vec{S}_{<k}$ the set of only those circle separations in $\vec{S}$ whose order is less than k.

Given fixed integers $m \geqslant 1$ and $n > 3$, we call a consistent orientation of S_k a *k-tangle* in S if it has no subset in

$$\mathcal{F} = \mathcal{F}_m^n := \{ F \subseteq 2^{\vec{U}} : \Big|\bigcap_{(A,B)\in F} B\Big| < m \text{ and } |F| < n \}.$$

A *tangle* in S is then a k-tangle for some k, and a *maximal tangle* in S is a tangle not contained in any other tangle in S. As usual, two tangles are *distinguishable* if neither of them is a subset of the other; a separation s *distinguishes* two tangles if they orient s differently, and s does so *efficiently* if it is of minimal order among all separations in S distinguishing that pair of tangles.

Using Lemma 16 we can show that there is a canonical nested set of circle separations which efficiently distinguishes all distinguishable tangles in S:

Theorem 18. *The set S of all circle separations of V contains a tree set $T = T(S)$ that efficiently distinguishes all distinguishable tangles of S. Moreover, this tree set T can be chosen canonically, i.e., so that for every automorphism α of S we have $T(S^\alpha) = T(S)^\alpha$.*

In order to prove Theorem 18 we need the following short lemma:

Lemma 5.15. *Let r and s be two circle separations of V. If r and s cross, then all four corner separations of r and s are again circle separations.*

Proof. Let $\vec{r} = (A, B)$ and $\vec{s} = (C, D)$. Since r and s cross, the sets $A \cap C$ and $B \cap D$ are non-empty and moreover intervals in the cyclic order. Thus $B \cup D$ is also an interval and therefore $\vec{r} \wedge \vec{s} = (A \cap C\,,\, B \cup D)$ is indeed a circle separation. □

Let us now prove Theorem 18.

Proof of Theorem 18. For every pair P, P' of distinguishable tangles in S let $\mathcal{A}_{P,P'}$ be the set of all circle separations that efficiently distinguish P and P'. We define a partial order $\preccurlyeq$ on the set of all pairs of distinguishable tangles by letting $\{\, P, P' \,\} \prec \{\, Q, Q' \,\}$ for two distinct such pairs if and only if the separations in $\mathcal{A}_{P,P'}$ have strictly lower order than those in $\mathcal{A}_{Q,Q'}$.

Let us show that the collection of these sets $\mathcal{A}_{P,P'}$ splinters hierarchically; the claim will then follow from Lemma 16.

For this let P, P' and Q, Q' be two distinguishable pairs of tangles in S and let $r \in \mathcal{A}_{P,P'}$ and $s \in \mathcal{A}_{Q,Q'}$. If r and s are nested, then r and s themselves are corner separations from different sides of r and s that lie in $\mathcal{A}_{P,P'}$ and $\mathcal{A}_{Q,Q'}$, respectively, in which case there is nothing to show.

So suppose that r and s cross. Then by Lemma 5.15 all corner separations of r and s are circle separations. By switching their roles if necessary we may assume that $|r| \leqslant |s|$; we shall treat the cases of $|r| < |s|$ and $|r| = |s|$ separately.

Let us first consider the case that $|r| < |s|$. Then $\{\, P, P' \,\} \prec \{\, Q, Q' \,\}$, so it suffices to show that (HS1) is satisfied, i.e., to find a corner separation of r and s in $\mathcal{A}_{Q,Q'}$. Since Q and Q' both orient s, which is of higher order than r, both Q and Q' also orient r. By $|r| < |s|$ and the efficiency of s, r cannot distinguish Q and Q'. Thus some orientation $\vec{r}$ of r lies in both Q and Q'.

By renaming them if necessary we may assume that $\vec{r} \in P$ and $\overleftarrow{r} \in P'$. Suppose now that one of $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \overleftarrow{s}$ has order at most $|s|$. Then Q and Q' would both orient that corner separation, and they would do so differently by the definition of a tangle. Thus that corner separation would lie in $\mathcal{A}_{Q,Q'}$, as desired.

Hence we may assume that both of $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \overleftarrow{s}$ have order higher than $|s|$. Then, by submodularity, both $\vec{r} \wedge \vec{s}$ and $\vec{r} \wedge \overleftarrow{s}$ have order less than $|r|$. Therefore both of these corner separations get oriented by P and P', but neither of them can distinguish P and P' by the efficiency of r. In fact by the consistency of P and P' we must have $(\vec{r} \wedge \vec{s}), (\vec{r} \wedge \overleftarrow{s}) \in P \cap P'$. However the set $\{\, \overleftarrow{r},\ (\vec{r} \wedge \vec{s}),\ (\vec{r} \wedge \overleftarrow{s}) \,\}$ lies in $\mathcal{F}$, contradicting the assumption that P and P' are tangles in S.

It remains to deal with the case that $|r| = |s|$ and show that (HS2) is satisfied. For this we shall find corner separations from different sides of r or of s that lie in $\mathcal{A}_{P,P'}$ and $\mathcal{A}_{Q,Q'}$, respectively. By the submodularity of the order function, and by switching the roles of r and s if necessary, we may assume that there are orientations of r and s such that both $\vec{r} \vee \vec{s}$ and $\vec{r} \vee \overleftarrow{s}$ have order at most $|r|$. By possibly renaming $\vec{s}$ and $\overleftarrow{s}$ we may further assume that $\vec{r} \vee \vec{s}$ distinguishes P and P'. Then, by the efficiency of r, we must have $|\vec{r} \vee \vec{s}| = |r|$, and hence $|\overleftarrow{r} \vee \overleftarrow{s}| \leqslant |s|$ by submodularity. Recall that we assumed $|\vec{r} \vee \overleftarrow{s}| = |r| = |s|$, so one of $\overleftarrow{r} \vee \overleftarrow{s}$ and $\vec{r} \vee \overleftarrow{s}$ must distinguish Q and Q'. Again, that corner separation must in fact distinguish Q and Q' efficiently, i.e., lie in $\mathcal{A}_{Q,Q'}$. Now this corner separation together with $\vec{r} \vee \vec{s}$ witnesses that (HS2) holds. $\square$

5.6. The splinter lemma for a relation

Arguably, the splinter lemma is not really concerned with a separation system itself any more but just with the nestedness relation. It can be phrased as a lemma about a reflexive, symmetric 'compatibility' relation. It then says that, if we have something like the fish lemma Lemma 2.1 to locally make our choice of elements from some collection $\mathfrak{A}$ more compatible, then we can make this compatibility global.

To make this precise, let A be some set and let $\sim$ be a reflexive and symmetric relation on A which we will refer to as '$\sim$-nested'. An element $c \in A$ is called a *$\sim$-corner of $a \in A$ and $b \in A$* if every element of A which is $\sim$-nested with both a and b is also $\sim$-nested with c. We now say that a collection $\mathfrak{A} = (\mathcal{A}_i)_{i\in I}$ of subsets of A *$\sim$-splinters* if for every pair of $\sim$-crossing (i.e., not $\sim$-nested) elements $a_i \in \mathcal{A}_i \setminus \mathcal{A}_j$ and $a_j \in \mathcal{A}_j \setminus \mathcal{A}_i$ there lies a $\sim$-corner of a_i and a_j in $\mathcal{A}_i \cup \mathcal{A}_j$.

Note that, under the general definition of a $\sim$-corner, both a and b always constitute $\sim$-corners of a and b. In the definition of '$\sim$-splinters' we can thus equivalently ask that for every pair of $\sim$-crossing elements $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$ there lies a $\sim$-corner of a_i and a_j in $\mathcal{A}_i \cup \mathcal{A}_j$. Note the omission of set-differences in the choice of a_i and a_j.

The splinter lemma and its proofs hold practically verbatim for this set of definitions:

Lemma 19 (Relation splinter lemma)**.** *Let A be a set, let $\mathfrak{A} = (\mathcal{A}_i)_{i\leqslant n}$ be a family of subsets of A and let $\sim$ be a reflexive and symmetric relation on A. If $\mathfrak{A}$ $\sim$-splinters, then we can pick an element a_i from each $\mathcal{A}_i$ so that $\{ a_1, \dots, a_n \}$ is pairwise $\sim$-nested.*

In a similar way one could generalize Lemma 16 to such a relation. However, we are forgoing this in favor of a version of Lemma 16 which not only works with a relation, but can also be used in the context of infinite families of infinite sets $\mathfrak{A}$. This will be the thin splinter lemma, Lemma 22.

5.7. Distinguishing directed tangles

To demonstrate the versatility of Lemma 19, we will now show an application to directed tangles of directed graphs. Directed tangles are an analogue to tangles in directed graphs, introduced by Giannopoulou, Kawarabayashi, Kreutzer, and Kwon [48]. They relate to the concept of directed tree-width proposed by Reed [63] and Johnson, Robertson, Seymour, and Thomas [56]. The 'directed separations' that directed tangles are defined with are more restrictive in terms of 'taking corners' than the usual separations of undirected graphs. This makes it impossible to find a tree of tangles, in the sense of a nested set of (directed) separations, for directed tangles. Instead Giannopoulou, Kawarabayashi, Kreutzer, and Kwon [48] introduced a weaker analogoue of trees of tangles for directed tangles: tangle-tree-labellings.

In this section we will establish an abstract version of directed separations, i.e., a directed version of abstract separation systems, and show how to retrieve an abstract version of such a tangle-tree-labelling using Lemma 19. We will use the results of [48] to guide our effort.

The key different of directed tangles to usual tangles of graphs is that they work with ‘directed separations’ of a digraph which have a ‘forwards’ and a ‘backwards’ direction. Formally, in the notation of [48], a *directed separation* (A, B) of a digraph G consists of vertex sets A and B with $A \cup B = V$, either such that no arc of G goes from $B \setminus A$ to $A \setminus B$ or such that no arc of G goes from $A \setminus B$ to $B \setminus A$. In the former case it is also written as $(A \to B)$ and in the latter case as $(A \leftarrow B)$. Note that every separation of the underlying undirected graph is a directed separation of G both ways, but generally G will have many more directed separations; the directed separations of G are a system of set separations of $V(G)$. The *order* of a directed separation is $|A \cap B|$, as usual, and we consider the usual partial order for set separations. A *directed k-tangle* in G or tangle of *order k,* then, is an orientation τ of all directed separations of order $< k$ such that for any three $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$ we have that $A_1 \cup A_2 \cup A_3 \neq V(G)$.

Diseparation systems

Differently from separations of a graph, the directed separations of a digraph do not form a universe of separations: for two directed separations $(A \to B)$ and $(C \to D)$ we have that $(A \cup C \to B \cap D)$ and $(A \cap C \to B \cup D)$ are also directed separations (called *upper* and *lower corner*, respectively). On the other hand, $(A \cup D, B \cap C)$ and $(A \cap D, B \cup C)$ generally may have cross-arcs going both ways, so the may not be directed separations.

To handle this, we will now develop an abstract version of the setting of directed separations.g We shall need to be able to keep track of the ‘forward’ direction of a directed separation, but also be able to speak of an orientation of a directed separation where we do not care about which direction is ‘forward’. The notation that follows, while a bit unwieldy, allows us to do so without clashing too much with existing notation of abstract separation systems – in fact, this will be a generalization of separation systems.

A *directed separation system*, or *diseparation system*, shall be a poset $\overleftrightarrow{S} = \vec{S} \cup \overleftarrow{S}$ together with an order-reversing involution $^*\colon \overleftrightarrow{S} \to \overleftrightarrow{S}$ which maps elements of $\vec{S}$ to elements of $\overleftarrow{S}$ and vice versa. Diseparations which are in $\vec{S}$, the *forwards diseparations*, are always denoted with a forwards arrow, those in $\overleftarrow{S}$, the *backwards diseparations*, with a backwards arrow. We extend the convention of writing *inverses* as $\overleftarrow{s} = \vec{s}^*$ and $\vec{s} = \overleftarrow{s}^*$. Diseparations which are only known to be in $\overleftrightarrow{S}$ are denoted with a double arrow, and * is used to denote their inverse. We note that every separation system can be turned into a diseparation system where $\vec{S} = \overleftarrow{S} = \overleftrightarrow{S}$. As with separation systems, we use s to denote the unoriented diseparation $\{\vec{s}, \overleftarrow{s}\}$ and $S = \{s : \vec{s} \in \vec{S}\}$.

If a diseparation system $\overleftrightarrow{U}$ is such that $\vec{U}$ (and thus also $\overleftarrow{U}$) is a lattice, we say that $\overleftrightarrow{U}$ is a *universe of directed separations*, or *diseparation universe* for short. This means that two diseparations $\vec{r}, \vec{s} \in \vec{U}$ have two *directed corners*: the *upper corner* $\vec{r} \vee \vec{s}$ and the *lower corner* $\overleftarrow{r} \vee \overleftarrow{s} = (\vec{r} \wedge \vec{s})^*$.[1] The directed separations of a digraph G form a diseparation universe as one would expect, and we denote this diseparation universe by $\overleftrightarrow{U}(G)$. The above translation of a separation system into a diseparation system also makes every univserse of separations a diseparation universe.

Given a diseparation system $\overleftrightarrow{S}$ an *orientation of* S is a set $O \subseteq \overleftrightarrow{S}$ such that, for every $s \in S$ we have that $|O \cap s| = 1$. Such an orientation O is *consistent* if there are no two distinct diseparations $\overleftrightarrow{s}, \overleftrightarrow{t} \in O$ such that $\overleftrightarrow{s}^* \leqslant \overleftrightarrow{t}$.

Given a diseparation universe $\overleftrightarrow{U}$, a diseparation system $\overleftrightarrow{S} \subseteq \overleftrightarrow{U}$ and some consistent orientation P of $\overleftrightarrow{S}$, we say that P is a *diprofile of* S if, for any $\vec{s}, \vec{t} \in \vec{S} \cap P$ we have that $(\vec{s} \vee \vec{t})^* \notin P$, and for any $\overleftarrow{s}, \overleftarrow{t} \in \overleftarrow{S} \cap P$ we have that $(\overleftarrow{s} \vee \overleftarrow{t})^* \notin P$.

A diseparation $\overleftrightarrow{s} \in \overleftrightarrow{S}$ is *small* if $\overleftrightarrow{s} \leqslant \overleftrightarrow{s}^*$, in which case $\overleftrightarrow{s}^*$ is *cosmall.* A diprofile P is *regular* if it contains no cosmall diseparation.

A function $|\cdot| : \overleftrightarrow{U} \to \mathbb{N}$ is called an *order function* if $|\vec{s}| = |\overleftarrow{s}|$ for all $\vec{s} \in \vec{U}$. An order function $|\cdot|$ called *submodular* if $|\vec{s}| + |\vec{t}| \geqslant |\vec{s} \vee \vec{t}| + |\vec{s} \wedge \vec{t}|$ for any two $\vec{s}, \vec{t} \in \vec{U}$, i.e., if it is submodular on the lattice $\vec{U}$; by symmetry $|\cdot|$ is then also submodular on $\overleftarrow{U}$.

Given such a submodular order function, the set $\overleftrightarrow{S}_k = \{\overleftrightarrow{s} \in \overleftrightarrow{U} : |\overleftrightarrow{s}| < k\}$ is a diseparation system. A diprofile of $\overleftrightarrow{S}_k$ is called a *k-diprofile in U* and a *diprofile in U* is a k-diprofile in U for some k. Note that directed k-tangles in a digraph G are regular k-diprofiles in the corresponding diseparation universe with the order function $|A, B| = |A \cap B|$.

5.7.1. Tangle tree-labellings

Giannopoulou, Kawarabayashi, Kreutzer, and Kwon's version of a tree-of-tangles theorem for directed tangles states the following:

Theorem 5.16 ([48, Theorem 6.2])**.** *Every set $\mathcal{T}$ of distinguishable tangles in a digraph G has a $\mathcal{T}$-tree-labelling.*

A directed separation (A, B) *distinguishes* two directed tangles τ, τ' if $(A, B) \in \tau$ and $(B, A) \in \tau'$ or vice versa. A pair of directed tangles is *distinguishable* if they are distinguished by some directed separation, and they are *ℓ-distinguishable* if the are distinguished by a directed separation of order at most ℓ, otherwise they are *ℓ-indistinguishable.*

The '$\mathcal{T}$-tree-labellings' that feature here may be seen as weakened versions of nested sets of separations as we know them from tree-of-tangles theorems, their

[1] Here we note that in diseparation universes, again, DeMorgans law's hold.

formalisation is closer to the 'trees of tangles' that we will discuss in Section 6.3. Formally, a *$\mathcal{T}$-tree-labelling* is a tuple (L, β, γ) consisting of a tree L, a bijection $\beta\colon V(L) \to \mathcal{T}$, and a map $\gamma\colon E(L) \to U(G)$ where

(D) for every path P in L, between vertices $x, y \in V(L)$ say, the minimum-order elements of $\{\, \gamma(e) : e \in P \,\}$ are efficient $\beta(x)$–$\beta(y)$-distinguishers.

Theorem 5.16 is proved first for sets of directed $> \ell$-tangles all of which induce the same directed ℓ-tangle:

Theorem 5.17 ([48, Theorem 6.3])**.** *Let $\mathcal{T}$ be a set of tangles of order $> \ell$ in a digraph G which are pairwise ℓ-distinguishable but $(< \ell)$-indistinguishable. Then there is a $\mathcal{T}$-tree-labelling (L, β, γ) of G such that $|\gamma(e)| = \ell$ for all $e \in E(L)$.*

With all the associated separations being of the same order, the condition (D) implies that every $\gamma(e)$ efficiently distinguishes all pairs of directed tangles across the two connected components of $L - e$, this implies that all the directed tangles in one component of $L - e$ contain the same orientation of $\gamma(e)$.

In fact, even [48]'s proof of Theorem 5.16 would allow replacing $\gamma\colon E(L) \to U(G)$ in the definition of a $\mathcal{T}$-tree-labelling by a function $\alpha\colon \vec{E}(L) \to \overleftrightarrow{U}(G)$ which commutes with the involution. We can then demand, in place of condition (D), the following condition:

(D') for every directed path P in L, from x to y say, the minimum-order elements of $\{\, \alpha(\vec{e}) : \vec{e} \in P \,\}$ are efficient $\beta(x)$–$\beta(y)$-distinguishers where $\alpha(\vec{e}) \in \beta(y)$ and $\alpha(\vec{e})^* \in \beta(x)$.

The existence of a $\mathcal{T}$-tree-labelling, as guaranteed by Theorem 5.17, i.e., for a fixed order ℓ, is then equivalent to the existence of a set of diseparations of G of order ℓ which induce a nested set of bipartitions of $\mathcal{P}$ under the map

$$\pi_{\mathcal{T}}\colon \overleftrightarrow{S}_{\ell+1}(G) \to \vec{\mathcal{B}}(\mathcal{T}),\ \overleftrightarrow{s} \mapsto (\{\, \tau \in \mathcal{T} : \overleftrightarrow{s}^* \in \tau \,\}, \{\, \tau \in \mathcal{T} : \overleftrightarrow{s} \in \tau \,\}),$$

which distinguishes all the elements of $\mathcal{P}$, since for such a nested set of bipartitions we can use the usual correspondence of nested set of separations to order-trees (cf. Theorem 2.2) to assert the existence of the tree L, and obtaining β is trivial.

We will now demonstrate how to prove the existence of such a set of diseprations using Lemma 19 in the abstract setting of diseparation systems and diprofiles. For that purpose, let us say that a diseparation s *distinguishes* two diprofiles P, P' if there exists an orientation $\overleftrightarrow{s}$ of s such that $\overleftrightarrow{s} \neq \overleftrightarrow{s}^*$ and $\overleftrightarrow{s} \in P$ but $\overleftrightarrow{s}^* \in P'$; it distinguishes them *efficiently* if it does so with lowest possible order. A pair of diprofiles in U is *ℓ-distinguishable* if they are distinguished by some diseparation of order at most ℓ. Two diprofiles in U are *distinguishable* if they are ℓ-distinguishable for some ℓ.

Given a set of diprofiles $\mathcal{P}$ of S, two diseparations $s, t \in S$ are called *$\mathcal{P}$-nested* if the bipartitions that they induce on the set $\mathcal{P}$ via $\pi_{\mathcal{P}} \colon \overleftrightarrow{S} \to \overrightarrow{\mathcal{B}}(\mathcal{P})$ form a nested set of separations, i.e., there is a pair of orientations of s and t which is contained in no diprofile in $\mathcal{P}$.

We observe that the diprofile property makes the map $\pi_{\mathcal{P}}$ compatible with meets and joins (i.e., $\pi_{\mathcal{P}}(\vec{s} \vee \vec{t}) = \pi_{\mathcal{P}}(\vec{s}) \vee \pi_{\mathcal{P}}(\vec{t})$ for all $\vec{s}, \vec{t} \in \overleftrightarrow{S}$) and with the involution. Thus, Lemma 2.1, the fish lemma, applies to the $\mathcal{P}$-nestedness-relation also, which ensures that directed corners of two disepartions are '$\mathcal{P}$-nestedness'-corners in the sense of Lemma 19.

Theorem 20. *Let $\overleftrightarrow{U}$ be a universe of directed separations with a submodular order function, let $\ell \in \mathbb{N}$ and let $\mathcal{P}$ be a set of diprofiles in U which are pairwise efficiently distinguished by diseparations of order precisely ℓ. Then there exists a $\mathcal{P}$-nested set of diseparations which distinguishes any two diprofiles in $\mathcal{P}$ efficiently.*

Proof. We let $\mathfrak{A} = (\mathcal{A}_{P,P'} : P, P' \in \mathcal{P},\ P \neq P')$ where

$$\mathcal{A}_{P,P'} = \{\, s \in S : s \text{ distinguishes } P \text{ and } P' \text{ efficiently} \,\},$$

and we will show that $\mathfrak{A}$ splinters with respect to the relation '$\mathcal{P}$-nested'.

Let $s \in \mathcal{A}_{P,P'}$ and $t \in \mathcal{A}_{Q,Q'}$ be two $\mathcal{P}$-crossing diseparations and note that both s and t have order ℓ. Since we could otherwise substitute s for t or vice versa, we may assume that P and P' both orient t the same, as $\overleftrightarrow{t}$ say, and that Q and Q' both orient s the same, as $\overleftrightarrow{s}$ say.

Let us first assume that $\overleftrightarrow{s}$ and $\overleftrightarrow{t}$ are both forwards or both backwards separations. Then by submodularity, one of $\overleftrightarrow{s} \vee \overleftrightarrow{t}$ and $\overleftrightarrow{s} \wedge \overleftrightarrow{t}$ is a diseparation of order at most ℓ.

If $\overleftrightarrow{s} \vee \overleftrightarrow{t}$ has order at most ℓ, then by the diprofile property and consistency it is in both $\mathcal{A}_{P,P'}$ and $\mathcal{A}_{Q,Q'}$. Moreover, $\overleftrightarrow{s} \vee \overleftrightarrow{t}$ is a '$\mathcal{P}$-nestedness'-corner of s and t, i.e., any diseparation which is $\mathcal{P}$-nested with s and t is also $\mathcal{P}$-nested with $\overleftrightarrow{s} \vee \overleftrightarrow{t}$, by our preliminary observation.

So suppose the order of $\overleftrightarrow{s} \vee \overleftrightarrow{t}$ is larger than ℓ, then the order of $\overleftrightarrow{s} \wedge \overleftrightarrow{t}$ is strictly less than ℓ. Since s and t are $\mathcal{P}$-crossing, there exists a diprofile $R \in \mathcal{P}$ with $\overleftrightarrow{s}^* \in R$ and $\overleftrightarrow{t}^* \in R$. The diprofile property implies that $(\overleftrightarrow{s} \wedge \overleftrightarrow{t})^* \in R$, but this means that R is $(< \ell)$-distinguishable from P, P', Q, and Q', each of which contradicts the choice of $\mathcal{P}$.

It remains to consider the case where $\overleftrightarrow{s}$ and $\overleftrightarrow{t}$ are one forwards and one backwards separation, say $\overleftrightarrow{s} = \vec{s} \in P'$ is forwards and $\overleftrightarrow{t} = \overleftarrow{t} \in Q$ is backwards. By submodularity, one of $\vec{s} \vee \vec{t}$ and $\vec{s} \wedge \vec{t}$ has order at most ℓ. If $\vec{s} \vee \vec{t}$ has order at most ℓ, then $\vec{s} \vee \vec{t} \in Q'$ by the diprofile property, and $(\vec{s} \vee \vec{t})^* \in Q$ by consistency. Together this implies $\vec{s} \vee \vec{t} \in \mathcal{A}_{Q,Q'}$. If $\vec{s} \wedge \vec{t}$ has order at most ℓ, then $(\vec{s} \wedge \vec{t})^* \in P$ by the diprofile property, and $\vec{s} \wedge \vec{t} \in P'$ by consistency. Together this implies $\vec{s} \vee \vec{t} \in \mathcal{A}_{P,P'}$. In either case the diseparation that we find is a '$\mathcal{P}$-nestedness'-corner of s and t.

Thus, we can apply Lemma 19 to obtain the desired $\mathcal{P}$-nested set. $\square$

Once we allow diseparations of different orders $\mathcal{P}$-nestedness becomes too strong of a condition to ask for: even if two '$\mathcal{P}$-crossing' diseparations are oriented by all the diprofiles in $\mathcal{P}$, we cannot necessarily uncross this pair as required for Lemma 19. $\mathcal{P}$-nested sets which distinguish all diprofiles in $\mathcal{P}$ may simply not exist, and Theorem 5.16 does not claim such a thing for directed tangles either. In fact, in a tangle-tree-labelling as constructed in the proof of Theorem 5.16 only the directed separations of lowest order are nested in this sense: if k is the lowest order of a directed separation in the $\mathcal{T}$-tree-labelling, then the directed separations of order k in this labelling are $\mathcal{T}_{k+1}$-nested, where $\mathcal{T}_{k+1}$ is the set of restrictions of the directed tangles in $\mathcal{T}$ to directed separations of order at most k.

One thing to note, when considering the separations of different order across a tangle-tree-labelling, is that one can construct multiple different tangle-tree-labellings using the same set of directed separations. This is reflected in the final part of the proof of Theorem 5.16 in [48] where one has to choose end-vertices for connecting two tree-labellings (obtained by induction) to each other by an edge.

There is, however, a weaker nestedness-like condition which the directed separations in the constructed tangle-tree-labellings *do* satisfy, and a construction as in the final part of the proof of Theorem 5.16 would allow us to again obtain a tangle-tree-labelling from a set of directed separations which is weakly nested in the following sense:

Given a set of diprofiles $\mathcal{P}$ in a diseparation universe U we say that two diseparations s and t are *weakly $\mathcal{P}$-nested* if

- if $|s| \leqslant |t|$, then there exists a *witnessing orientation* $\overleftrightarrow{s}$ *of* s such that of any two diprofiles $P, Q \in \mathcal{P}$ which are $(< |s|)$-indistinguishable and are distinguished by t at least one contains $\overleftrightarrow{s}$; and
- if $|s| \geqslant |t|$, then there exists a *witnessing orientation* $\overleftrightarrow{t}$ *of* t such that of any two diprofiles $P, Q \in \mathcal{P}$ which are $(< |t|)$-indistinguishable and are distinguished by s at least one contains $\overleftrightarrow{t}$.

A set of separations is *weakly $\mathcal{P}$-nested* if every pair of diseparations from that set is weakly $\mathcal{P}$-nested. To help understand the equality case, we establish the following lemma.

Lemma 5.18. *If $|s| = |t| = \ell$ and s and t are not weakly $\mathcal{P}$-nested, then there are four diprofiles $P_{\vec{s},\vec{t}}, P_{\vec{s},\overleftarrow{t}}, P_{\overleftarrow{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}} \in \mathcal{P}$ which are pairwise $(< \ell)$-indistinguishable where $\overleftrightarrow{s}, \overleftrightarrow{t} \in P_{\overleftrightarrow{s},\overleftrightarrow{t}}$ for $\overleftrightarrow{s} \in \{\vec{s}, \overleftarrow{s}\}, \overleftrightarrow{t} \in \{\vec{t}, \overleftarrow{t}\}$.*

Proof. As s and t are not weakly $\mathcal{P}$-nested, we may suppose without loss of generality that there is a pair of diprofiles $P_{\vec{s},\vec{t}}, P_{\vec{s},\overleftarrow{t}} \in \mathcal{P}$ which is $(< |s|)$-indistinguishable (with $\vec{s} \in P_{\vec{s},\vec{t}}, P_{\vec{s},\overleftarrow{t}}$, $\vec{t} \in P_{\vec{s},\vec{t}}$ and $\overleftarrow{t} \in P_{\vec{s},\overleftarrow{t}}$) and that there is a pair of diprofiles $P_{\overleftarrow{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}} \in \mathcal{P}$ which is $(< |s|)$-indistinguishable (with $\overleftarrow{s} \in P_{\overleftarrow{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}}$, $\vec{t} \in P_{\overleftarrow{s},\vec{t}}$ and $\overleftarrow{t} \in P_{\overleftarrow{s},\overleftarrow{t}}$).

If not all four of these are pairwise $(< |s|)$-indistinguishable, there is a diseparation r such that $|r| < |s|$ and $\vec{r} \in P_{\vec{s},\vec{t}} \cap P_{\vec{s},\overleftarrow{t}}$ and $\overleftarrow{r} \in P_{\overleftarrow{s},\vec{t}} \cap P_{\overleftarrow{s},\overleftarrow{t}}$.

However, then there is a pair of directed corners between $\vec{r}$ and $\vec{t}$, and one of these two has order less than $|t|$ and thereby contradicts the choice of one of the two pairs of diprofiles. $\square$

For the weak $\mathcal{P}$-nestedness-relation we show only a very weak version of Lemma 2.1, which is just enough to show that the directed corners we will find when trying to apply Lemma 19 are 'weak $\mathcal{P}$-nestedness'-corners.

Lemma 5.19. *Let s and t be two diseparations which are not weakly $\mathcal{P}$-nested, let r be weakly $\mathcal{P}$-nested with both s and t and let c be a directed corner of s and t. If $|s| \leqslant |t|$ and c efficiently distinguishes a pair of diprofiles in $\mathcal{P}$ which are also efficiently distinguished by t but not by s, then r is weakly $\mathcal{P}$-nested with c.*

Proof. If $|c| \leqslant |r|$ we need to show that there is a witnessing orientation of c for weak $\mathcal{P}$-nestedness with r. In this case $|s| \leqslant |t| = |c| \leqslant |r|$, so there are orientations $\overleftrightarrow{s}$ and $\overleftrightarrow{t}$ of s and t witnessing the weak $\mathcal{P}$-nestedness with r.

If c has an orientation $\overleftrightarrow{c}$ such that $\overleftrightarrow{c} \leqslant \overleftrightarrow{s}$ or $\overleftrightarrow{c} \leqslant \overleftrightarrow{t}$, then that orientation of c is suitable to witness weak $\mathcal{P}$-nestedness with r. Otherwise, $\overleftrightarrow{c}^* = \overleftrightarrow{s}^* \vee \overleftrightarrow{t}^*$ is an orientation of c, and by symmetry we can consider just the case that $\overleftrightarrow{c}$ is an upper corner, i.e., $\overleftrightarrow{c} = \vec{c} = \vec{s} \vee \vec{t}$ which means that $\overleftrightarrow{s} = \vec{s}$ and $\overleftrightarrow{t} = \vec{t}$.

If some $P, Q \in \mathcal{P}$ which are $(< |t|)$-indistinguishable and distinguished by r, say $\overleftrightarrow{r} \in P$ and $\overleftrightarrow{r}^* \in Q$, are such that $\overleftarrow{c} \in P \cap Q$, then by diprofile property we have $\overleftarrow{s} \in P$ or $\overleftarrow{t} \in P$, and $\overleftarrow{s} \in Q$ or $\overleftarrow{t} \in Q$. Moreover, by the $\mathcal{P}$-nestedness of r with s and t, we may suppose up to renaming that $\overleftarrow{s} \in P$ and $\overleftarrow{t} \in Q$, but $\vec{s} \in P$ and $\vec{t} \in Q$. This implies that $|s| = |t|$ since P and Q are $(< |t)|$-indistinguishable.

Consider the diprofiles $P_{\vec{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}}$ from Lemma 5.18. We claim that $P, Q, P_{\vec{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}}$ are pairwise $(< |s|)$-indistinguishable. Otherwise, there is a separation u such that $|u| < |s|$ and $\vec{u} \in P \cap Q$, $\overleftarrow{u} \in P_{\vec{s},\vec{t}} \cap P_{\overleftarrow{s},\overleftarrow{t}}$, but then there is a directed corner between u and s of order less than $|s|$ which distinguishes P from Q or $P_{\vec{s},\vec{t}}$ from $P_{\overleftarrow{s},\overleftarrow{t}}$, a contradiction.

Now either $\vec{r}$ or $\overleftarrow{r}$ is in $P_{\overleftarrow{s},\overleftarrow{t}}$. In the first case $P_{\overleftarrow{s},\overleftarrow{t}}$ and Q contradict the assumption that $\vec{t}$ witnesses the weak $\mathcal{P}$-nestedness of t with r, whereas in the second case $P_{\overleftarrow{s},\overleftarrow{t}}$ and P contradict the assumption that $\vec{s}$ witnesses the weak $\mathcal{P}$-nestedness of s with r. This concludes the proof, that there exists a witnessing orientation of c if $|c| \leqslant |r|$.

If we have $|r| \leqslant |c|$, then we need to show that r has a witnessing orientation for weak $\mathcal{P}$-nestedness with c. Suppose for a contradiction that r has no such orientation.

We first want to show, that we can assume $|r| \leqslant |s|$. So, suppose $|s| < |r| \leqslant |t| = |c|$, then there is an orientation $\overleftrightarrow{r}$ of r witnessing the weak $\mathcal{P}$-nestedness with t. Since this orientation of r does not witness also to the weak $\mathcal{P}$-nestedness with c, there is a pair $P, Q \in \mathcal{P}$ of $(< |r|)$-indistinguishable diprofiles that are distinguished by c where both contain $\overleftrightarrow{r}^*$. By the weak $\mathcal{P}$-nestedness of r and t and the fact that

P, Q both contain $\overleftrightarrow{r}^*$, we know that this pair is not distinguished by t. They are then, by the diprofile property, distinguished by s and thus $|s| = |r|$, by efficiency, which is a contradiction.

We may thus assume $|r| \leqslant |s| \leqslant |t| = |c|$. Then, there are orientations of r which witness the weak $\mathcal{P}$-nestedness with s and t, respectively. If these two orientations coincide, this orientation clearly witness also the weak $\mathcal{P}$-nestedness with c, so suppose that $\overleftrightarrow{r}$ witnesses the weak $\mathcal{P}$-nestedness with s and $\overleftrightarrow{r}^*$ witnesses the weak $\mathcal{P}$-nestedness with t.

By the premise of this lemma, c and t efficiently distinguish a pair of diprofiles $P_{\vec{t},\overleftrightarrow{s}}, P_{\overleftarrow{t},\overleftrightarrow{s}} \in \mathcal{P}$ which is not distinguished by s, so $\overleftrightarrow{s} \in P_{\vec{t},\overleftrightarrow{s}}, P_{\overleftarrow{t},\overleftrightarrow{s}}$ and $\vec{t} \in P_{\vec{t},\overleftrightarrow{s}}, \overleftarrow{t} \in P_{\overleftarrow{t},\overleftrightarrow{s}}$. Since s and t are not weakly $\mathcal{P}$-nested, there is also a pair $P_{\vec{t},\overleftrightarrow{s}^*}, P_{\overleftarrow{t},\overleftrightarrow{s}^*}$ of $(< |s|)$-indistinguishable diprofiles in $\mathcal{P}$ which is distinguished by t and which both contain $\overleftrightarrow{s}^*$.

We claim that $P_{\vec{t},\overleftrightarrow{s}}, P_{\overleftarrow{t},\overleftrightarrow{s}}, P_{\vec{t},\overleftrightarrow{s}^*}$ and $P_{\overleftarrow{t},\overleftrightarrow{s}^*}$ are pairwise $(< |s|)$-indistinguishable. Indeed, if u is a diseparation of order $< |s|$ such that $\overleftrightarrow{u} \in P_{\vec{t},\overleftrightarrow{s}} \cap P_{\overleftarrow{t},\overleftrightarrow{s}}$ and $\overleftrightarrow{u}^* \in P_{\vec{t},\overleftrightarrow{s}^*} \cap P_{\overleftarrow{t},\overleftrightarrow{s}^*}$, there is a directed corner of u and t which either contradicts the fact that $P_{\vec{t},\overleftrightarrow{s}}$ and $P_{\overleftarrow{t},\overleftrightarrow{s}}$ are $(< |t|)$-indistinguishable or the fact that $P_{\vec{t},\overleftrightarrow{s}^*}$ and $P_{\overleftarrow{t},\overleftrightarrow{s}^*}$ are $(< |s|)$-indistinguishable.

So, since $P_{\vec{t},\overleftrightarrow{s}}$ and $P_{\overleftarrow{t},\overleftrightarrow{s}^*}$ are distinguished by both s and t and the witnessing orientations of r for s and t are distinct, $P_{\vec{t},\overleftrightarrow{s}}$ and $P_{\overleftarrow{t},\overleftrightarrow{s}^*}$ need also be distinguished by r. Similarly, $P_{\vec{t},\overleftrightarrow{s}^*}$ and $P_{\overleftarrow{t},\overleftrightarrow{s}}$ are also distinguished by r. Thus, there either is an orientation $\overleftrightarrow{r}$ of r such that $\overleftrightarrow{r} \in P_{\vec{t},\overleftrightarrow{s}} \cap P_{\overleftarrow{t},\overleftrightarrow{s}}$ and $\overleftrightarrow{r}^* \in P_{\vec{t},\overleftrightarrow{s}^*} \cap P_{\overleftarrow{t},\overleftrightarrow{s}^*}$, or there is an orientation $\overleftrightarrow{r}$ of r such that $\overleftrightarrow{r} \in P_{\vec{t},\overleftrightarrow{s}} \cap P_{\vec{t},\overleftrightarrow{s}^*}$ and $\overleftrightarrow{r}^* \in P_{\overleftarrow{t},\overleftrightarrow{s}} \cap P_{\overleftarrow{t},\overleftrightarrow{s}^*}$.

In the former case, the two diprofiles $P_{\vec{t},\overleftrightarrow{s}}$ and $P_{\overleftarrow{t},\overleftrightarrow{s}}$ are distinguished by t and $(< |r|)$-indistinguishable (since $|r| \leqslant |s|$). Since both diprofiles contain $\overleftrightarrow{r}$ and r is weakly $\mathcal{P}$-nested with t, the orientation $\overleftrightarrow{r}$ must be the orientation witnessing weak $\mathcal{P}$-nestedness with t. But $P_{\vec{t},\overleftrightarrow{s}^*}$ and $P_{\overleftarrow{t},\overleftrightarrow{s}^*}$ are also distinguished by t and $(< |r|)$-indistinguishable and both contain $\overleftrightarrow{r}^*$ which contradicts the conclusion that $\overleftrightarrow{r}$ was the witnessing orientation.

The latter case contradicts the weak $\mathcal{P}$-nestedness of r and s by a symmetrical argument. □

We are now in a position to prove our main theorem with the help of Lemma 19.

Theorem 21. *For every set $\mathcal{P}$ of distinguishable diprofiles in a diseparation universe with a submodular order function, there exists a weakly $\mathcal{P}$-nested set of separations which efficiently distinguishes all pairs of diprofiles in $\mathcal{P}$.*

Proof. We define $\mathfrak{A} = (\mathcal{A}_{P,P'} : P, P' \in \mathcal{P}, P \neq P')$ as before and want to apply Lemma 19 with the relation 'weakly $\mathcal{P}$-nested'.

Now, in order to show that we can apply Lemma 19, let $s \in \mathcal{A}_{P,P'} \setminus \mathcal{A}_{Q,Q'}$ and $t \in \mathcal{A}_{Q,Q'} \setminus \mathcal{A}_{P,P'}$ not be weakly $\mathcal{P}$-nested and suppose that $|s| \leqslant |t|$. Let $\overleftrightarrow{s}$ be the orientation of s with $\overleftrightarrow{s} \in Q \cap Q'$.

Let us first consider the case that $|s| < |t|$. Then there needs to be, since s and t are not weakly $\mathcal{P}$-nested, a pair $R, R' \in \mathcal{P}$ of diprofiles which are $(< |s|)$-indistinguishable such that $\overleftrightarrow{s}^* \in R \cap R'$, $\vec{t} \in R$ and $\overleftarrow{t} \in R'$. Now if $|\overleftrightarrow{s} \vee \vec{t}| \leqslant |t|$, it is in $\mathcal{A}_{Q,Q'}$ and a 'weak $\mathcal{P}$-nestedness'-corner by Lemma 5.19. Otherwise $|\overleftrightarrow{s} \wedge \vec{t}| < |s|$, by submodularity, which contradicts the fact that R, R' are $(< |s|)$-indistinguishable.

In the case that $|s| = |t|$, either there are diprofiles $R, R' \in \mathcal{P}$ as above – $(< |s|)$-indistinguishable and with $\overleftrightarrow{s}^* \in R \cap R'$, $\vec{t} \in R$, and $\overleftarrow{t} \in R'$ – in which case the same argument applies, or we can exchange the roles of R and R', i.e., there are $(< |s|)$-indistinguishable diprofiles $R, R' \in \mathcal{P}$ with $\overleftrightarrow{t}^* \in R \cap R'$, $\vec{s} \in R$, and $\overleftarrow{s} \in R'$, where $\overleftrightarrow{t}$ is the orientation of t in $P \cap P'$ – in which case a symmetrical argument holds.

This shows that we can apply Lemma 19, so there exists a weakly $\mathcal{P}$-nested set of diseparations which efficiently distinguishes all the diprofiles in $\mathcal{P}$. □

We note that in proving this Lemma, we did not need to rely on the concept of $\mathcal{P}$-nestedness that we established for Theorem 20, in particular not on Theorem 20 itself.

One can turn a weakly $\mathcal{P}$-nested set N of diseparations which efficiently distinguishes $\mathcal{P}$ as obtained from Theorem 21 into the equivalent of a tangle-tree-labelling following [48]'s proof of Theorem 5.16. We will not go into the full proof here, but only sketch the construction.

We call a tree T together with a map $\alpha \colon \vec{E}(T) \to \overleftrightarrow{S}$ an *S-tree* if it commutes with the involution. (This is a natural extension of the notion of S-trees from separation systems.) Our aim is to find an S-tree (T, α) where the vertices of T are diprofiles in $\mathcal{P}$ and which satisfies condition (D') (where β is the identity function).

Assume that N is a $\subseteq$-minimal such distinguishing set and let T be a tree with vertex set $\mathcal{P}$ where we draw an edge for every $s \in N$ between some arbitrary pair of diprofiles in $\mathcal{P}$ which in N is efficiently distinguished by s and s only.

We show that T is a tree, inductively, by showing that every contraction $T_k = T/\sim_k$, where $P \sim_k Q$ if P and Q are $(< k)$-indistinguishable, is a tree. Let $\mathcal{P}_k$ be the set of all k-diprofiles induced by diprofiles in $\mathcal{P}$, we identify the vertices of T_k with the according diprofiles in $\mathcal{P}_k$.

For the induction start, observe that all edges in T_1 are labelled with diseparations of order 0. The definition of weak $\mathcal{P}$-nestedness ensures that these are $\mathcal{P}_1$-nested, i.e., the induced bipartitions of $\mathcal{P}_1$ under $\pi_{\mathcal{P}_1}$ are nested. Pulling back the tree we obtain for these bipartitions from the equivalence of tree-sets and tree-orders (e.g. Theorem 2.2), we obtain an S-tree which satisfies the condition that every diseparation is the only one in N efficiently distinguishing the diprofiles at its end-vertices. Moreover, any two diprofiles which are not adjacent in this tree are efficiently distinguished by at least two diseparations in N, and thus do not share an edge in T_1. Thus, T_1 is uniquely defined: it is isomorphic to that S-tree.

For the induction step, we only need to show that for every $P \in \mathcal{P}_k$ the induced subgraph on the set $\mathcal{P}_P \subseteq \mathcal{P}_{k+1}$ of those $(k+1)$-diprofiles which induce P is a tree. Then, by the induction hypothesis we know that T_k is a tree and thus these induced subgraphs are joined up in a tree-like way, making T_{k+1} a tree.

Every $s \in N$ which efficiently distinguishes two diprofiles in $\mathcal{P}_P$ efficiently distinguishes $k+1$-profiles only in $\mathcal{P}_P$: if s were to efficiently distinguish some diprofiles $Q, Q' \in \mathcal{P}_{k+1} \setminus \mathcal{P}_P$, say, then N contains some efficient P–Q-distinguisher, which would not be weakly $\mathcal{P}$-nested with s.

Now, the diseparations in N which efficiently distinguish $\mathcal{P}_P$ induce a nested set of bipartitions of $\mathcal{P}_P$, i.e., they are $\mathcal{P}_P$-nested. Then, again, the induced subgraph on $\mathcal{P}_P$ is uniquely determined to be the S-tree corresponding to these bipartitions.

By induction, this shows that T is indeed a tree. We then chooce as α the function which assigns to each edge the separation for which we drew that edge in the construction of T, with the orientations corresponding to profiles at the end-vertices. The contractions T_k can then also be used to verify that (D') is satisfied: if some diprofiles Q and Q' are efficiently distinguished at order ℓ, say, then Q and Q' live in the same node of $T_{\ell-1}$, associated to the common restriction of Q and Q' to order $\ell - 1$ which we call P. The Q–Q'-path in T thus uses only diseparations associated with edges of order at least ℓ. However, Q and Q' are contained in distinct nodes in T_ℓ, so the lowest order along the Q–Q'-path is ℓ. Moreover, any diseparations of order ℓ along the path distinguish some pair of profiles in $\mathcal{P}_P$ efficiently, and we know that these are $\mathcal{P}_P$-nested. This ensures (D').

5.8. The thin splinter lemma: A canonical splinter lemma for infinite settings

Lemma 13 is proved by induction: it finds a separation $a_i \in \mathcal{A}_i$ which is nested with some element of every other $\mathcal{A}_j$, and then proceeds inductively on the remaining $n-1$ family members, restricted to those separations nested with a_i. This approach cannot deal with infinite families of sets, however.

In this section we overcome these difficulties and present a way to obtain a version of Lemma 13 for infinite families of sets of separations. Combining ideas from the canonical splinter lemma, Lemma 16, and the relation splinter lemma, Lemma 19, we obtain the 'thin' splinter lemma, which implies the existing tree-of-tangle theorems for infinite graphs (Theorem 5.20 and 5.21 below).

⇝ p. 95 **Lemma 22** (Thin splinter lemma)**.** *If* $(\, \mathcal{A}_i : i \in I\,)$ *thinly splinters with respect to some reflexive symmetric relation* $\sim$ *on* $\mathcal{A} := \bigcup_{i\in I} \mathcal{A}_i$*, then there is a set* $N \subseteq \mathcal{A}$ *which meets every* $\mathcal{A}_i$ *and is nested, i.e.,* $n_1 \sim n_2$ *for all* $n_1, n_2 \in N$*. Moreover, this set* N *can be chosen invariant under isomorphisms: if* φ *is an isomorphism between* $(\mathcal{A}, \sim)$ *and* $(\mathcal{A}', \sim')$*, then we have* $N((\,\varphi(\mathcal{A}_i) : i \in I\,)) = \varphi(N((\,\mathcal{A}_i : i \in I\,)))$.

We prove this statement in Section 5.8.1. Like Lemma 13, the statement of this theorem is a bit technical, as we want it to be as widely applicable as possible.

We will see applications of the thin splinter lemma in Section 5.9.3.

5.8.1. Statement and proof

Our approach to generalizing the finite splinter lemma into an infinite setting will require that the separations involved do not, in a sense, cross too badly in that they cross only finitely many separations of lower order. This will allow us to choose separations that minimise the number of separations crossing them, an idea which also appeared in Carmesin's original proof of Theorem 5.20 in [10], as well as in [14] and our proof of the canonical splinter lemma, Lemma 16. However, our theorem here applies to a more general setting and will allow us directly to deduce Carmesin's theorem for locally finite graphs.

In order to also be able to deduce the full Theorem 5.20 for arbitrary graphs, we will state our theorem in more generality here: not as a theorem about nestedness and separations, but as a theorem about a general nestedness-like relation, like in Section 5.6. This allows us to apply the theorem in Section 5.9.3 not to separations directly, where it would fail, but to substitute separators as a proxy giving our Theorem 24. From this result we will retrieve the separations for our proof of Theorem 5.20 in Section 5.9.3, but we will also build from this a tree of tree-decompositions to deduce Theorem 5.21 in Section 5.9.3.

The statement of our Lemma 22 is also inspired by our canonical splinter lemma for the finite setting in [38], and it too will result in a *canonical* nested set, a set which is invariant under isomorphisms.

So let $\mathcal{A}$ be some set and $\sim$ a reflexive and symmetric binary relation on $\mathcal{A}$. In analogy to our terminology for separation systems, we say that two elements a and b of $\mathcal{A}$ are *nested* if $a \sim b$. Elements of $\mathcal{A}$ that are not nested *cross*. As usual, a subset of $\mathcal{A}$ is nested if all of its elements are pairwise nested, and a single element is nested with a set N if it is nested with every element of N.

In an abuse of notation, given elements a and b of $\mathcal{A}$, we call $c \in \mathcal{A}$ a *corner* of a and b if every element of $\mathcal{A}$ crossing c also crosses one of a and b. Observe that with this definition corners of elements of $\mathcal{A}$ exhibit the same behaviour as was asserted by Lemma 2.1 for corner separations. However, in contrast to the terminology of separation systems, we do not insist here that a corner of a and b is itself nested with both a and b. This distinction will become relevant in Section 5.9.3.

Now let $(\mathcal{A}_i : i \in I)$ be a family of non-empty subsets of $\mathcal{A}$ and $|\cdot| : I \to \mathbb{N}_0$ some function, where I is a possibly infinite index set. We shall think of $|i|$ as the order of the elements of $\mathcal{A}_i$. For an $a \in \mathcal{A}$ and $k \in \mathbb{N}_0$ the *k-crossing number* of a is the number of elements of $\mathcal{A}$ that cross a and lie in some $\mathcal{A}_i$ with $|i| = k$. This k-crossing number is either a natural number or infinity. The family $(\mathcal{A}_i : i \in I)$ *thinly splinters* if it satisfies the following three conditions:

(TS1) For every $i \in I$ all elements of $\mathcal{A}_i$ have finite k-crossing number for all $k \leqslant |i|$.

(TS2) If $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$ cross with $|i| < |j|$, then $\mathcal{A}_j$ contains some corner of a_i and a_j that is nested with a_i.

(TS3) If $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$ cross with $|i| = |j| = k \in \mathbb{N}_0$, then either $\mathcal{A}_i$ contains a corner of a_i and a_j with strictly lower k-crossing number than a_i, or else $\mathcal{A}_j$ contains a corner of a_i and a_j with strictly lower k-crossing number than a_j.

We are now ready to state and prove the main result of this section:

Lemma 22 (Thin splinter lemma)**.** *If $(\mathcal{A}_i : i \in I)$ thinly splinters with respect to some reflexive symmetric relation $\sim$ on $\mathcal{A} := \bigcup_{i\in I} \mathcal{A}_i$, then there is a set $N \subseteq \mathcal{A}$ which meets every $\mathcal{A}_i$ and is nested, i.e., $n_1 \sim n_2$ for all $n_1, n_2 \in N$. Moreover, this set N can be chosen invariant under isomorphisms: if φ is an isomorphism between $(\mathcal{A},\sim)$ and $(\mathcal{A}',\sim')$, then we have $N((\varphi(\mathcal{A}_i) : i \in I)) = \varphi(N((\mathcal{A}_i : i \in I)))$.*

Proof. We shall construct inductively, for each $k \in \mathbb{N}_0$, a nested set $N_k \subseteq \mathcal{A}$ extending N_{k-1} and meeting every $\mathcal{A}_i$ with $|i| \leqslant k$, so that the choice of N_k is invariant under isomorphisms. The desired nested set N will then be the union of all these sets N_k.

We set $N_{-1} := \emptyset$. Suppose that for some $k \in \mathbb{N}_0$ we have already constructed a nested set N_{k-1} so that N_{k-1} is canonical and meets every $\mathcal{A}_i$ with $|i| \leqslant k-1$. We shall construct a canonical nested set $N_k \supseteq N_{k-1}$ that meets every $\mathcal{A}_i$ with $|i| \leqslant k$.

Let N_k^+ be the set consisting of the following: for every $i \in I$ with $|i| = k$, among those elements of $\mathcal{A}_i$ that are nested with N_{k-1}, those of minimum k-crossing number. We claim that $N_k \coloneqq N_{k-1} \cup N_k^+$ is as desired.

Since the choice of N_k^+ is invariant under isomorphisms, and N_{k-1} is canonical by assumption, N_k is clearly canonical as well. It thus remains to show that N_k meets every $\mathcal{A}_i$ with $|i| = k$, and that the set N_k is nested.

To see that the former is true, let $i \in I$ with $|i| = k$ be given. It suffices to show that $\mathcal{A}_i$ contains some element that is nested with N_{k-1}. If $\mathcal{A}_i$ already meets N_{k-1} there is nothing to show, so suppose that it does not. By condition (TS1) every element of $\mathcal{A}_i$ crosses only finitely many elements of N_{k-1}; pick an $a_i \in \mathcal{A}_i$ that crosses as few as possible. Suppose for a contradiction that a_i crosses some element of N_{k-1}, that is, some $a_j \in \mathcal{A}_j$ with $|j| < |i|$. But then, by condition (TS2), $\mathcal{A}_i$ contains a corner of a_i and a_j that is nested with a_j. This element of $\mathcal{A}_i$ does not cross a_j and therefore, by virtue of being a corner of a_i and a_j, crosses fewer elements of N_{k-1} than a_i does, contrary to the choice of a_i. Therefore, N_k indeed contains an element of each $\mathcal{A}_i$ with $|i| \leqslant k$.

Let us now show that N_k is nested. Since N_{k-1} is a nested set by assumption, and every element of N_k^+ is nested with N_{k-1}, we only need to show that the set N_k^+ itself is nested. So suppose that some two elements of N_k^+ cross. These two elements then are some $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$ with $|i| = |j| = k$. But now condition (TS3) asserts that one of $\mathcal{A}_i$ and $\mathcal{A}_j$ contains a corner of a_i and a_j with a strictly lower k-crossing number than the corresponding element a_i or a_j. Since both a_i and a_j are nested with N_{k-1} their corner is nested with N_{k-1} as well, and hence contradicts the choice of a_i or a_j for N_k^+. □

5.9. Distinguishing tangles in infinite graphs

In this section we are going to apply Lemma 22 to vertex tangles (or more generally: robust regular profiles) in infinite graphs. We will later show another application of Lemma 22, not to (vertex) tangles of graphs, though still to infinite graphs, using it to exhibit the edge-connectivity of a graph in Section 5.10.

As a simple example, we start with applying it to tangles in locally finite graphs in Section 5.9.2. This application is a straightforward application to a universe of separations and demonstrates the prototypical way of applying Lemma 22.

It is also possible to apply Lemma 22 to arbitrary infinite graphs, and we do so in Section 5.9.3. This application uses another new, and interesting, shift of perspective: We cannot apply Lemma 22 directly to the sets of separations efficiently distinguishing two profiles since, in general, these will not splinter thinly. Instead, we consider the separators of those separations, there Lemma 22 does apply:

Theorem 24. *Given a set of distinguishable robust regular profiles $\mathcal{P}$ of a graph G there exists a canonical nested set of separators efficiently distinguishing any pair of profiles in $\mathcal{P}$.* ⇝ p. 109

This theorem acts as an intermediary result between the existing results about tangles in arbitrary infinite graphs. On the one hand we can, waiving canonicity, transform the nested set of separators back into a nested set of separations, recovering the following result of Carmesin about distinguishing tangles in infinite graphs by way of a nested set of separations:

Theorem 5.20 ([10, Theorem 5.12])**.** *For any graph G, there is a nested set N of separations that distinguishes efficiently any two robust principal profiles (that are not restrictions of one another).*

This theorem is a cornerstone in Carmesin's proof that every infinite graph has a tree-decomposition displaying all its topological ends. For more about the relation between ends and tangles also see [18, 57]. We deduce Theorem 5.20 from Theorem 24 in Section 5.9.3.

On the other hand, if we want to keep canonicity, we can use Theorem 24 to deduce a result by Carmesin, Hamann, and Miraftab [14]. They construct a canonical object, which they call a tree of tree-decompositions, to distinguish the tangles:

Theorem 5.21 ([14, Remark 8.3])**.** *Let G be a connected graph and $\mathcal{P}$ a distinguishable set of principal robust profiles in G. There exists a canonical tree of tree-decompositions with the following properties:*

(1) the tree of tree-decompositions distinguishes $\mathcal{P}$ efficiently;
(2) if $t \in V(T)$ has level k, then $(T_t, \mathcal{V}_t)$ contains only separations of order k;
(3) nodes t at all levels have $|V(T_t)|$ neighbours on the next level and the graphs assigned to them are all torsos of $(T_t, \mathcal{V}_t)$.

We will deduce Theorem 5.21 from Theorem 24 in Section 5.9.3, but Theorem 24 is also an interesting result in its own right: the set of separators that it provides is a natural intermediate object between the non-canonical nested set of separations in Theorem 5.20 and the canonical tree of tree-decompositions in Theorem 5.21.

Moreover, proving Theorem 5.20 or Theorem 5.21 by first proving Theorem 24 and then deducing them breaks up the proof nicely and is, in total, shorter than the original proofs from [10, 14].

5.9.1. Terminology and basic facts

Recall that a profile P in G is *regular* if it does not contain any cosmall separation of G, i.e., it contains no separation of the form $(V(G), X)$. Note that, in graphs, the irregular profiles are not of large interest, since they always point towards either the empty set or a single non-cut-vertex. Formally, we can summarize this statement from [23] as follows:

Lemma 5.22 ([23])**.** *Let $G = (V, E)$ be a graph and P an irregular profile in G then either G is connected and $P = \{ (V, \emptyset) \}$ or G has a non-cutvertex $x \in V$ such that*

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

These irregular profiles are distinguished efficiently from each other and from all other profiles in G by the set of separations

$$\{ \{V(G), \emptyset\} \} \cup \{ \{V(G), \{x\}\} : x \in V(G) \text{ and } x \text{ is not a cutvertex of } G \}.$$

Every separation in this set is nested with all separations of G. Hence, our efforts for applications in graphs will concentrate on regular profiles.

Given some set of vertices $X \subseteq V(G)$, we say that a connected component C of $G - X$ is *tight*, if $N(C) = X$.

For two vertices $x, y \in V(G)$ of a graph G, an *x–y-separator of order k* is a vertex set $X \subseteq V(G) \setminus \{ x, y \}$ of size k such that x and y lie in different components of $G - X$. We shall need the following basic fact about such separators in infinite graphs at various points throughout our applications.

Lemma 5.23 ([54, 2.4])**.** *Let G be a graph, $u, v \in V(G)$ and $k \in \mathbb{N}$. Then there are only finitely many separators of size at most k separating u and v minimally.*

Additionally, we shall use the following more general observation about separations nested with a corner separation:

Lemma 5.24. *Let $\vec{r}$ and $\vec{s}$ be two separations. Every separation nested with one of r or s is also nested with at least one of $\vec{r} \wedge \vec{s}$ and $\vec{r} \vee \vec{s}$.*

Proof. Let t be a separation nested with, say, r. Then t has an orientation $\vec{t}$ with either $\vec{t} \leqslant \vec{r}$ or $\vec{t} \leqslant \overleftarrow{r}$. In the first case t is nested with $\vec{r} \vee \vec{s}$ by $\vec{t} \leqslant \vec{r} \leqslant (\vec{r} \vee \vec{s})$. In the latter case t is nested with $\vec{r} \wedge \vec{s}$ by $\vec{t} \leqslant \overleftarrow{r} \leqslant (\vec{r} \wedge \vec{s})^*$. $\square$

The separations that distinguish a given pair of profiles exhibit a lattice structure:

Lemma 5.25. *Let $\vec{U}$ be a universe with a submodular order function and P and P' two profiles in $\vec{U}$. If $\vec{r}, \vec{s} \in P$ distinguish P and P' efficiently, then both $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ also lie in P and distinguish P and P' efficiently.*

Proof. If one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ has order at most $|r| = |s|$, then that corner separation lies in P and distinguishes P and P' by their consistency and the profile property (P). The efficiency of r and s now implies that neither of the two considered corner separations can have order strictly lower than $|r|$. Therefore, by submodularity, both of them have order exactly $|r|$, which implies the claim. $\square$

We will utilize the fact that separations which efficiently distinguish two regular profiles are 'tight'. For a set $X \subseteq V$ a component C of $G - X$ is *tight* if $N(C) = X$. We say that a separation (A, B) of G is *tight* if for $X := A \cap B$ each of $A \setminus B$ and $B \setminus A$ contains some tight component of $G - X$.

Lemma 5.26. *Let P, P' be two distinct regular profiles in an arbitrary graph G. If (A, B) is a separation of finite order that efficiently distinguishes P and P', then (A, B) is tight.*

Proof. Let $(A, B) \in P$, $(B, A) \in P'$.

Suppose for a contradiction that $B \setminus A$ does not contain a tight component of $G - (A \cap B)$. Let $Y_1, \dots Y_m$ be an enumeration of all proper subsets of $A \cap B$. For every Y_l let $\mathcal{C}_l$ be the set of components of $G - (A \cap B)$ in B with neighbourhood exactly Y_l. By consistency of P' we have $(\bigcup \mathcal{C}_l \cup Y_l, V \setminus \bigcup \mathcal{C}_l) \in P'$. Since moreover (A, B) efficiently distinguishes P from P' and $|Y_l| < |A \cap B|$ we know that $(\bigcup \mathcal{C}_l \cup Y_l, V \setminus \bigcup \mathcal{C}_l) \in P$ as well. Moreover, $(A \cap B, V) \in P$ since P is regular. Thus, by an inductive application of the profile property (P) we have that for every l

$$(A \cap B, V) \vee (\bigcup \mathcal{C}_1 \cup Y_1, V \setminus \bigcup \mathcal{C}_1) \vee \dots \vee (\bigcup \mathcal{C}_l \cup Y_l, V \setminus \bigcup \mathcal{C}_l) \in P.$$

However, for $l = m$ this contradicts the assumption since

$$(A \cap B, V) \vee (\bigcup \mathcal{C}_1 \cup Y_1, V \setminus \bigcup \mathcal{C}_1) \vee \dots \vee (\bigcup \mathcal{C}_m \cup Y_m, V \setminus \bigcup \mathcal{C}_m) = (B, A) \notin P\,.$$

$\square$

Moreover, we shall need a way to transition between nested sets of separations and tree-decompositions of graphs. Such a method already exists in finite graphs [13], and Kneip and Gollin [49] all but proved an analogue for infinite graphs. We shall combine Theorem 2.2 (which is from [49]) with the ingredients of the proof in [13] to show Lemma 5.27.

Recall that a tree-decomposition of a G is a pair $(T, \mathcal{V})$ of a tree T together with a family $\mathcal{V} = (V_t)_{t \in T}$ of vertex sets $V_t \subseteq V(G)$ such that:

(T1) $V(G) = \bigcup_{t \in T} V_t$;

(T2) Given $e \in E[G]$ there exists a $t \in T$ such that $e \subseteq V_t$;

(T3) Given a path P in T from t_1 to t_3 and a vertex $t_2 \in P$ we have $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$.

A separation (A, B) is induced by a tree-decomposition $(T, \mathcal{V})$ if, and only if, there exists an edge $tt' \in T$ such that for the components $T_t, T_{t'}$ of $T - tt'$ containing t or t' respectively, we have

$$(A, B) = \left(\bigcup_{t'' \in T_t} V_{t''}, \bigcup_{t'' \in T_{t'}} V_{t''} \right).$$

Also recall that Theorem 2.2 states that every regular tree set which does not contain a chain of order type $\omega + 1$ is isomorphic to the edge tree set of a suitable tree.

We obtain the following lemma, whose proof is inspired by [13].

Lemma 5.27. *Let $G = (V, E)$ be an infinite graph and let $N \subseteq S_{\aleph_0}(G)$ be a regular tree set. If we have for every ω-chain $(A_1, B_1) < (A_2, B_2) < \ldots$ which is contained in N that $\bigcap_{i\in\mathbb{N}} B_i = \emptyset$, then there exists a tree-decomposition $(T, \mathcal{V})$ of G whose set of induced separations is $\vec{N}$.*

Moreover, this tree-decomposition can be chosen canonical: if $\varphi\colon G \to G'$ is an isomorphism of graphs, then the tree-decomposition constructed for $\varphi(N)$ in G' is precisely the image under φ of the tree-decomposition constructed for N in G.

Proof. Let $T = (V, E)$ be the tree obtained from applying Theorem 2.2 to N. Note that, by Theorem 2.2, any isomorphism between the edge tree sets of any two trees induces an isomorphism of the underlying trees.

Let α be the isomorphism from the edge tree set of T to N. Given some node $t \in T$ let us denote as F_t the set of oriented separations

$$F_t := \{\, \alpha(s,t) : (s,t) \in \vec{E} \,\}\,.$$

We define the bags of our tree-decomposition as $V_t := \bigcap_{(A,B)\in F_t} B$. Let us verify that $(T, \mathcal{V})$ with $\mathcal{V} = (\, V_t \,)_{t\in T}$ is the desired tree-decomposition.

For (T1) let $v \in V$ be given; we need to find a $t \in T$ with $v \in V_t$. If $v \in A \cap B$ for some $(A, B) \in \vec{N}$ then $v \in V_t$ for t being either of the two end-vertices of the edge whose image under α is (A, B). Otherwise, v induces an orientation O of $E(T)$ by orienting each edge $\{\, x, y \,\}$ of T as (x, y) if $v \in B \setminus A$ for $(A, B) = \alpha(x, y)$.

Observe that O is consistent. If O has a sink, that is, if there is a node t of T all of whose incident edges are oriented inwards by O, then $v \in V_t$ by definition of O. If O does not have a sink then O contains an ω-chain. This is impossible though, since by definition of O we would have $v \in \bigcap_{i\in\mathbb{N}} B_i$, where (A_i, B_i) is the image under α of the i-th element of the ω-chain in O. Thus (T1) holds.

The proof of (T2) can be carried out in much the same way due to the fact that every edge of G is included in either A or B for each $(A, B) \in \vec{N}$.

Before we prove (T3), let us show that $(T, \mathcal{V})$ indeed induces N. For this we need to show that if (x, y) is an oriented edge of T then

$$\alpha(x,y) = \left(\bigcup_{z\in T_x} V_z \,,\, \bigcup_{z\in T_y} V_z \right) ,$$

where T_x and T_y are the components of $T - xy$ containing x and y, respectively. So let $(x, y) \in \vec{E}$ be given and $\alpha(x, y) = (A, B)$. Observe first that $A \cap B \subseteq V_x \cap V_y$ by definition. It thus suffices to show that $A \supseteq \bigcup_{z\in T_x} V_z$ and $B \supseteq \bigcup_{z\in T_y} V_z$ to establish the desired equality.

To see this, consider a vertex $v \in V_z$ for some $z \in T_x$. Let $\vec{e}$ be the first edge of the unique z–x-path in T and let $\alpha(\vec{e}) = (A', B')$. We have $\vec{e} \leqslant (x, y)$ by definition of an edge tree set, and hence $(A', B') \leqslant (A, B)$ since α is an isomorphism. From this we know that $A' \subseteq A$. We further have $(B', A') \in F_z$ and thus, by definition of V_z, that $v \in A'$. This shows $v \in A$. The argument that $B \supseteq \bigcup_{z \in T_y} V_z$ is similar.

Having established that $(T, \mathcal{V})$ indeed induces N, we can now deduce from this that (T3) holds: if V_{t_1} and V_{t_3} are two bags of $(T, \mathcal{V})$ which both contain some vertex v, then v also needs to lie in the separator of every separation that is an image under α of an edge on the path P in T from t_1 to t_3. Therefore, v lies in every V_{t_2} with $t_2 \in P$. $\square$

5.9.2. Profiles in locally finite graphs

In this section we apply Lemma 22 to the set of separations of a locally finite graph, which will result in a canonical nested set of separations efficiently distinguishing any two distinguishable regular profiles in G. The proof of this theorem will be a straightforward application of Lemma 22 to sets $\mathcal{A}_{P,P'}$ of separations efficiently distinguishing two profiles in G. Following the strategy of this proof, one might be able to obtain similar results for other infinite separation systems, e.g., in a matroid.

So let $G = (V, E)$ be a locally finite connected graph and $\mathcal{P}$ a set of robust regular profiles in G.

Let I be the set of pairs of distinguishable profiles in $\mathcal{P}$. For each pair P and P' of distinguishable profiles in $\mathcal{P}$ let $\mathcal{A}_{P,P'}$ be the set of all separations of G that distinguish P and P' efficiently. Observe that by definition all separations in $\mathcal{A}_{P,P'}$ are of the same order; let us write $|P, P'|$ for this order.

Let $\mathcal{A}$ be the union of all the $\mathcal{A}_{P,P'}$. We wish to show that $(\,\mathcal{A}_i : i \in I\,)$ thinly splinters, using as the relation $\sim$ on $\mathcal{A}$ the usual nestedness of separations. We shall prove first that condition (TS1) is satisfied, i.e., that each separation in an $\mathcal{A}_{P,P'}$ crosses only finitely many other separations from sets $\mathcal{A}_{Q,Q'}$ with $|Q, Q'| \leqslant |P, P'|$.

Making use of the tightness of the separations in the $\mathcal{A}_{P,P'}$, condition (TS1) will follow immediately from the following assertion:

Proposition 5.28. *Let (A, B) be a separation that efficiently distinguishes some two regular profiles in G. Then G has only finitely many tight separations of order at most $|(A, B)|$ that cross (A, B).*

We shall derive Proposition 5.28 from the following lemma about tight separations:

Lemma 5.29. *Let (A, B) and (A', B') be two tight separations of G. Then (A', B') is either nested with (A, B), or its separator $A' \cap B'$ is a $\subseteq$-minimal x–y-separator in G for some pair x, y of vertices from $(A \cap B) \cup N(A \cap B)$.*

Proof. Since (A', B') is tight each of $A' \setminus B'$ and $B' \setminus A'$ contains some tight component of $G - (A' \cap B')$. If $A \cap B$ meets all tight components of $G - (A' \cap B')$

then in particular $A \cap B$ meets these two components, say in x and in y. But then $A' \cap B'$ is a $\subseteq$-minimal x–y-separator with $x, y \in A \cap B$.

Therefore, we may assume that $A \cap B$ misses some tight component C' of $G - (A' \cap B')$. By switching their names if necessary we may assume that this component C' is contained in $A \setminus B$. Since $C' \subseteq A \setminus B$ has no neighbours in $B \setminus A$ but has $A' \cap B'$ as its neighbourhood we can infer that $(A' \cap B') \subseteq A$.

Consider now a tight component C of $G - (A \cap B)$ that is contained in $B \setminus A$. From $(A' \cap B') \subseteq A$ it follows that C does not meet $A' \cap B'$ and is hence contained in either $A' \setminus B'$ or $B' \setminus A'$. By possibly switching the roles of A' and B' we may assume that $C \subseteq A' \setminus B'$. As above we can conclude from the tightness of C that $(A \cap B) \subseteq C$.

It remains to check two cases. If $(B \setminus A) \cap (B' \setminus A')$ is empty we have $B \subseteq A'$ and $B' \subseteq A$, that is, that (A', B') is nested with (A, B). The other remaining case is that $(B \setminus A) \cap (B' \setminus A')$ is non-empty.

In that case, since G is connected, the set $(A \cap B) \cap (A' \cap B')$ must be non-empty as well, since $N((B \setminus A) \cap (B' \setminus A')) \subseteq (A \cap B) \cap (A' \cap B')$. Pick a vertex z from that set. Since (A', B') is tight z has neighbours x and y in some tight components of $G - (A' \cap B')$ contained in $A' \setminus B'$ and in $B' \setminus A'$, respectively. Then $A' \cap B'$ is a $\subseteq$-minimal x–y-separator in G, and moreover $x, y \in (A \cap B) \cup N(A \cap B)$ since $z \in A \cap B$. □

Let us now use Lemma 5.29 to establish Proposition 5.28:

Proof of Proposition 5.28. Since G is locally finite the set $(A \cap B) \cup N(A \cap B)$ is finite. Therefore, by Lemma 5.23, there are only finitely many $\subseteq$-minimal x–y-separators of size at most $|(A, B)|$ with $x, y \in (A \cap B) \cup N(A \cap B)$. Leveraging again the fact that G is locally finite, and using that G is connected, we get that there are only finitely many separations of G with such a separator.

The assertion now follows from Lemma 5.29 since we know by Lemma 5.26 that (A, B) is tight. □

The family $(\mathcal{A}_i : i \in I)$ therefore satisifes condition (TS1). We show condition (TS2) in the following lemma:

Lemma 5.30. *Let $(A, B) \in \mathcal{A}_{P,P'}$ and $(C, D) \in \mathcal{A}_{Q,Q'}$ with $|(A, B)| < |(C, D)|$. Then some corner separation of (A, B) and (C, D) lies in $\mathcal{A}_{Q,Q'}$.*

In other words, if $\vec{a}_i \in \mathcal{A}_i$ and $\vec{a}_j \in \mathcal{A}_j$ cross with $|i| < |j|$, then $\mathcal{A}_j$ contains some corner separation of a_i and a_j.

Proof. Since $|(A, B)| < |(C, D)|$ it follows that both Q and Q' orient $\{A, B\}$ the same, say $(A, B) \in Q \cap Q'$. If $|(A, B) \vee (C, D)| \leqslant |(C, D)|$ or $|(A, B) \vee (D, C)| \leqslant |(C, D)|$, it follows that this corner separation efficiently distinguishes Q and Q' by Lemma 2.1, so suppose that this is not the case. Then submodularity implies that $|(B, A) \vee (C, D)| < |(A, B)|$ and $|(B, A) \vee (D, C)| < |(A, B)|$, which in turn contradicts the efficiency of (A,B), since one of $(B,A) \vee (C,D)$ and $(B,A) \vee (D,C)$ would also distinguish the two robust profiles P and P'. $\square$

It remains to show that $(\mathcal{A}_i : i \in I)$ satisfies condition (TS3). As a first step we consider separations $(A,B) \in \mathcal{A}_{P,P'}$ and $(C,D) \in \mathcal{A}_{Q,Q'}$ of equal order, and show that there are two opposite corner separations of (A,B) and (C,D) that lie in $\mathcal{A}_{P,P'}$ or in $\mathcal{A}_{Q,Q'}$:

Lemma 5.31. *Let $(A,B) \in \mathcal{A}_{P,P'}$ and $(C,D) \in \mathcal{A}_{Q,Q'}$ with $|(A,B)| = |(C,D)|$. Then there is either a pair of two opposite corner separations of (A,B) and (C,D) with one element in $\mathcal{A}_{P,P'}$ and one in $\mathcal{A}_{Q,Q'}$, or else there are two pairs of opposite corner separations of (A,B) and (C,D), the first with both elements in $\mathcal{A}_{P,P'}$ and the second with both elements in $\mathcal{A}_{Q,Q'}$.*

Proof. From $|(A,B)| = |(C,D)|$ it follows that P and P' both orient $\{C,D\}$, and likewise that Q and Q' both orient $\{A,B\}$.

Let us first treat the case that one of P and P' orients both $\{A,B\}$ and $\{C,D\}$ in the same way as one of Q and Q' does. So suppose that, say, both P and Q contain (A,B) as well as (C,D).

If P' contains (D,C), then (C,D) is in $\mathcal{A}_{P,P'}$ and Lemma 5.25 gives that $(A,B) \vee (C,D) \in \mathcal{A}_{P,P'}$ and $(B,A) \vee (D,C) \in \mathcal{A}_{P,P'}$. Thus by property (P) we also have $(A,B) \vee (C,D) \in \mathcal{A}_{Q,Q'}$, producing the desired pair of opposite corner separations. If Q' contains (B,A) we argue analogously.

So suppose that $(C,D) \in P'$ and $(A,B) \in Q'$. Then $(B,A) \vee (C,D) \in P'$ and $(A,B) \vee (D,C) \in Q'$ by the profile property, since by submodularity and the efficiency of (A,B) and (C,D) both of these corner separations have order exactly $|(A,B)|$. These two separations, then, are opposite corner separations of (A,B) and (C,D) with the first lying in $\mathcal{A}_{P,P'}$ and the second lying in $\mathcal{A}_{Q,Q'}$.

The remaining case is that no two of the four profiles agree in their orientation of $\{A,B\}$ and $\{C,D\}$. But then both of (A,B) and (C,D) lie in $\mathcal{A}_{P,P'}$ as well as in $\mathcal{A}_{Q,Q'}$, and the existence of two pairs of opposite corner separations, one with both elements in $\mathcal{A}_{P,P'}$ and one with both in $\mathcal{A}_{Q,Q'}$, follows from Lemma 5.25 and the disagreement of the four profiles on $\{A,B\}$ and $\{C,D\}$. $\square$

We note that Lemma 5.30 and 5.31 do not require the graph to be locally finite, and we will be reusing these lemmas in Section 5.9.3. For now, let us show that $(\mathcal{A}_i : i \in I)$ satisfies condition (TS3) using Lemma 5.24 and 5.31:

Lemma 5.32. *If $\vec{a}_i \in \mathcal{A}_i$ and $\vec{a}_j \in \mathcal{A}_j$ cross with $k = |i| = |j|$, then either $\mathcal{A}_i$ contains a corner separation of a_i and a_j with strictly lower k-crossing number than $\vec{a}_i$, or else $\mathcal{A}_j$ contains a corner separation of a_i and a_j with strictly lower k-crossing number than $\vec{a}_j$.*

Proof. By switching their roles if necessary we may assume that the k-crossing number of $\vec{a}_i$ is at most the k-crossing number of $\vec{a}_j$.

From Lemma 5.31 it follows that $\mathcal{A}_j$ contains a corner separation of a_i and a_j whose opposite corner separation lies in either $\mathcal{A}_i$ or $\mathcal{A}_j$. Now Lemma 5.24 implies that the sum of the k-crossing numbers of this pair of opposite corner separations is at most the sum of the k-crossing numbers of $\vec{a}_i$ and $\vec{a}_j$. This inequality is in fact strict since $\vec{a}_i$ and $\vec{a}_j$ cross each other but are each nested with both corner separations.

If the first corner separation is not already as desired, that is, if its k-crossing number is not strictly lower than the k-crossing number of $\vec{a}_j$, we can infer that the k-crossing number of the opposite corner separation is strictly lower than that of $\vec{a}_i$. Since we assumed in the beginning that the k-crossing number of $\vec{a}_i$ is no greater than that of $\vec{a}_j$ this proves the claim. □

This all but completes the proof the main result of this subsection, which is similar to [14, Theorem 7.5]:

Theorem 23. *Let G be a locally finite connected graph and $\mathcal{P}$ some set of robust regular profiles in G. Then there exists a nested set $\mathcal{N}$ of separations which efficiently distinguishes any two distinguishable profiles in $\mathcal{P}$. Moreover, this set is canonical, i. e., invariant under isomorphisms: if $\alpha\colon G \to G'$ is an isomorphism, then $\alpha(\mathcal{N}(G, \mathcal{P})) = \mathcal{N}(\alpha(G), \alpha(P))$.*

Proof. The combination of Proposition 5.28 and Lemma 5.30 and 5.32 shows that the family $(\,\mathcal{A}_i : i \in I\,)$ thinly splinters. The nested set $N \subseteq \mathcal{A}$ produced by Lemma 22 meets each set $\mathcal{A}_i$ and thus disinguishes all pairs of dinstinguishable profiles in $\mathcal{P}$ efficiently. □

The nested set found by Theorem 23 does not in general correspond to a tree-decomposition of G: even in locally finite graphs it is not generally possible to find a *tree-decomposition* which efficiently distinguishes all the distinguishable robust regular bounded profiles, as witnessed by the following example:

Example 5.33. Consider the graph displayed in Fig. 5.4. This graph is constructed as follows: for every $n \in \mathbb{N}$ pick a copy of $K^{2^{n+2}}$ together with $n + 3$ vertices $w_1^n, \dots, w_{n+3}^n$. Pick 2^n vertices of the $K^{2^{n+2}}$ and call them $u_1^n, \dots, u_{2^n}^n$. Additionally, pick 2^{n+1} vertices from $K^{2^{n+2}}$, disjoint from the set of u_i^n, and call them $v_1^n, \dots, v_{2^{n+1}}^n$. Now identify u_i^{n+1} with v_i^n and add edges between every w_i^n and every w_j^{n+1} as well as between w_i^n and $v_1^n = u_1^{n+1}$.

Finally, we pick one copy of K^{10} and join one vertex v_1^0 of this K^{10} to u_1^1 and u_1^2. Additionally we pick two vertices w_1^0, w_2^0 which are distinct from v_1^0 from this K^{10} and add an edge between each w_i^0 and each w_j^1.

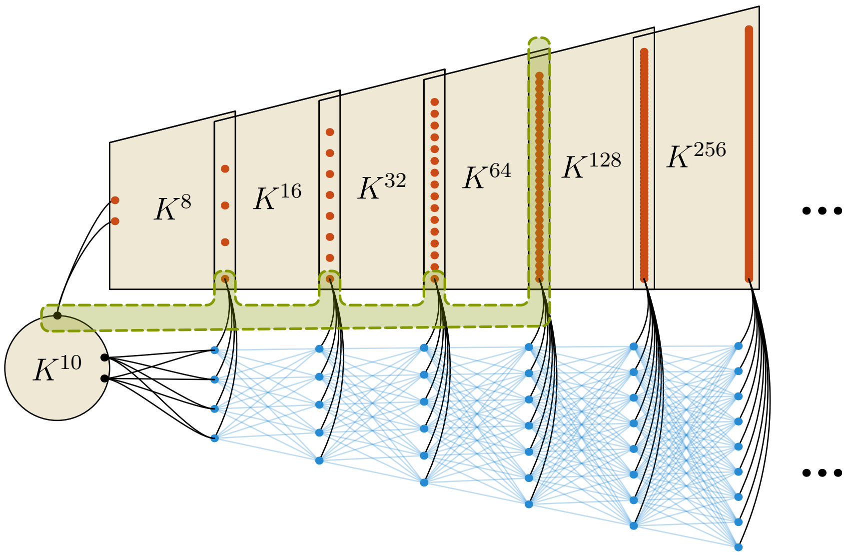

Figure 5.3.: A locally finite graph where no tree-decomposition distinguishes all the robust regular bounded profiles efficiently. The green separator is the one of the only separation which efficiently distinguishes the profile induced by the K^{64} from the profile induced by the K^{128}.

Now each of the chosen $K^{2^{n+2}}$ induces a robust profile P_n of order $\frac{2}{3} \cdot 2^{n+1}$ which obviously is regular and bounded. The only separation which efficiently distinguishes P_n and P_{n+1} is the separation s_n with separator $\{ v_1^i : i < n \} \cup \{u_j^{n+1}\}$.

Additionally, the K^{10} induces a robust profile P_0 of order 4. However, the only separation that efficiently distinguishes P_0 and P_1 has the separator $\{ v_1^0, w_1^0, w_2^0 \}$. But these separations $s_1, s_2, \dots$, and s_0 can be oriented such as to form a chain of order type $\omega + 1$. This chain witnesses that there cannot be a tree-decomposition which distinguishes all regular bounded profiles efficiently: the separations given by such a tree-decomposition would have to contain this chain of order type $\omega + 1$ which is not possible as every chain in the edge tree set of a tree has length at most ω, cf. Lemma 5.27.

Theorem 23 can however be used to show that for every fixed integer k the subset of $\mathcal{N}$ consisting of all separations of order at most k gives rise to a tree-decomposition of G, as this subset will satisfy the conditions from Lemma 5.27. In particular we can use Theorem 23 together with Lemma 5.27 to prove [14, Theorem 7.3] that there is, for every $k \in \mathbb{N}$, every locally finite graph G and every set $\mathcal{P}$ of robust regular profiles which are pairwise distinguishable by a separation of order at most k, a canonical tree-decomposition of G that efficiently distinguishes all profiles from $\mathcal{P}$.

5.9.3. Profiles in graphs with vertices of infinite degree

When we consider graphs with vertices of infinite degree, the method of the previous section fails as we loose Proposition 5.28: It does not necessarily hold that every separation in an $\mathcal{A}_{P,P'}$ crosses only finitely many other separations from sets $\mathcal{A}_{Q,Q'}$ with $|Q,Q'| \leqslant |P,P'|$. Moreover, Dunwoody and Krön [30] gave an example of a graph which does not contain a canonical nested set of separations separating its ends. As ends induce robust regular profiles, in arbitrary graphs, it is not generally possible to find a canonical nested set of separations distinguishing all the robust regular profiles.

To show the result for locally finite graphs we made use of the observation that only finitely many distinct *separators* are involved, and then used that every separator appears in only finitely many separations. Thus, in this section instead of applying Lemma 22 directly to some set of *separations*, we are going to apply it to only the set of *separators*.

With this approach we show that in an arbitrary graph you can find a canonical nested set of separators which efficiently distinguishes all the robust regular profiles in G. We shall make the meaning of this more precise shortly. We propose that this set of separators is a natural intermediate object for distinguishing profiles. Moreover, we will show that if we restrict ourselves to the set of robust principal profiles – which we will define at the end of this section – then from this set we can build both a non-canonical nested set of separations as in Theorem 5.20 (from [10]) as well as a canonical tree of tree-decompositions in the sense of [14].

Either of these objects can trivially be converted back to a set of separators. Our technique splits the process of building either of these cleanly into two independent steps, which makes it more accessible than the proofs in [10] and [14]. Moreover, the first step of this process also works for non-principal but regular profiles, allowing us to also get a (intermediate) result for those profiles, unlike the theorems from [10] and [14]. Note that distinguishing non-principal profiles is also discussed extensively in [43].

Many of the techniques applied throughout are similar to or inspired by arguments made in [14], particularly the approach of minimising the crossing-number, even though the different levels of abstraction make it hard to draw concrete parallels.

Let us now begin with the formal notation. We say that *a set of vertices* $X \subseteq V(G)$ *efficiently distinguishes* a pair P and P' of profiles in G if there exists a separation (A, B) of G with separator $A \cap B = X$ which efficiently distinguishes P and P'. Such a separation (A, B) is then a *witness* that X efficiently distinguishes P and P'.

Given some set of distinguishable robust regular profiles $\mathcal{P}$ of an (infinite) graph G, we define as $\mathcal{A}$ the set of all such separators X which distinguish some pair of profiles in $\mathcal{P}$ efficiently. We say that a separator X is *nested* with $Y \in \mathcal{A}$, i.e., $X \sim Y$, whenever X is contained in $C \cup Y$ for some component C of $G - Y$. In other words Y does not properly separate any two vertices of X. This relation is reflexive, the following lemma shows that it is also symmetric on $\mathcal{A}$. Unfortunately, its natural extension to all finite subsets of $V(G)$ is not. The reader should take note that this will lead to some situations where we argue that some set Y is nested with some $X \in \mathcal{A}$ *provided that* $Y \in \mathcal{A}$.

Lemma 5.34. *If $X, Y \in \mathcal{A}$ and X is contained in Y together with some component of $G - Y$, then Y is contained in X together with some component of $G - X$.*

Proof. Pick a separation (A, B) witnessing that $X \in \mathcal{A}$. Since this separation efficiently distinguishes two regular profiles, by Lemma 5.26, there are at least two tight components of $G - X$, one in either side of (A, B). At least one of these tight components, say C, does not meet Y and is therefore contained in a connected component C' of $G - Y$. Now, as required, we find

$$X = N(C) \subseteq C \cup N(C) \subseteq C' \cup N(C') \subseteq C' \cup Y. \qquad \square$$

As usual, we take as I the set of pairs of distinguishable profiles in $\mathcal{P}$. But this time we define $\mathcal{A}_{P,P'}$ for each pair P, P' in I to be the set of all the sets of vertices in G which distinguish P and P' efficiently. All these separators in $\mathcal{A}_{P,P'}$ have the same size; this size shall be $|P, P'|$.

We claim that $\{ \mathcal{A}_{P,P'} : \{P, P'\} \in I \}$ thinly splinters. Before we can show condition (TS1) we need to make two basic observations about how the vertices of a crossing pair of separators in $\mathcal{A}$ lie:

Lemma 5.35. *If $X, Y \in \mathcal{A}$ cross, then Y contains a vertex from every tight component of $G - X$.*

Proof. If C is a tight component of $G - X$ such that Y does not contain any vertex of C then C is contained in some component C' of $G - Y$. However, then $X = N(C) \subseteq C' \cup Y$, i.e., X is nested with Y contradicting the assertion. $\square$

Lemma 5.36. *If $X, Y \in \mathcal{A}$ cross, then Y contains a pair of vertices v and w such that X is a $\subseteq$-minimal v–w-separator.*

Proof. There are at least two tight components C_1, C_2 of $G - X$ and Y meets both of them by Lemma 5.35. Let v be a vertex in $Y \cap C_1$ and w a vertex in $Y \cap C_2$. As both C_1 an C_2 are tight components, X is indeed a $\subseteq$-minimal v–w-separator. $\square$

We can now combine these with Lemma 5.23 to prove condition (TS1).

Lemma 5.37. *For every pair of profiles $P, P' \in \mathcal{P}$ every $X \in \mathcal{A}_{P,P'}$ has finite k-crossing-number for all $k \leqslant |P, P'|$.*

Proof. By Lemma 5.36, for every $Y \in \mathcal{A}$ of size k which crosses X, there are vertices $v, w \in X$ which are minimally separated by Y. However, there is only a finite number of pairs of vertices v, w in X and by Lemma 5.23 every pair has only finitely many minimal separators of size k. Therefore, only finitely many such $Y \in \mathcal{A}$ exist. $\square$

The following lemmas show how the separators of corner separations behave under our new nestedness relation. We will need these to prove conditions (TS2) and (TS3). Recall from Section 5.8.1 that a *corner* of two separators $X, Y \in \mathcal{A}$ is a separator $Z \in \mathcal{A}$ which crosses only elements of $\mathcal{A}$ which cross either X or Y. Note that this does not imply that Z is nested with X and Y.

Lemma 5.38. *Let $X, Y \in \mathcal{A}$ be a crossing pair of separators and let (A_X, B_X) and (A_Y, B_Y), respectively, be separations which witness that these are in $\mathcal{A}$. Then for every $Z \in \mathcal{A}$ which is nested with both X and Y there is a component C_Z of $G - Z$, such that $X \cup Y \subseteq C_Z \cup Z$. In particular, $(A_X \cup A_Y) \cap (B_X \cap B_Y)$, the separator of $(A_X, B_X) \vee (A_Y, B_Y)$, is a corner of X and Y provided that it lies in $\mathcal{A}$.*

Proof. We first show that Z does not separate X and Y. Since Z is nested with X and X efficiently distinguishes two regular profiles there is, by Lemma 5.26, a tight component C_X of $G - X$ which is disjoint from Z. By Lemma 5.35, there is a vertex $y \in C_X \cap Y \subseteq Y \setminus Z$.

By a symmetrical argument there also exists a vertex $x \in X \setminus Z$. Since C_X is tight there is a path from x to y contained in C_X except for x. This path avoids Z.

Now, since Z is nested with X there is a component C_Z of $G - Z$ which contains $X \setminus Z$. In particular this component contains x. Similarly, there is a component of $G - Z$ containing $Y \setminus Z$ and hence, in particular, y. Since Z does not separate x and y this component is the same as C_Z. Therefore, $X \cup Y \subseteq C_Z \cup Z$, as required. In particular, if $(A_X \cup A_Y) \cap (B_X \cap B_Y) \in \mathcal{A}$ then $(A_X \cup A_Y) \cap (B_X \cap B_Y) \subseteq C_Z \cup Z$, hence $(A_X \cup A_Y) \cap (B_X \cap B_Y) \sim Z$ and therefore $(A_X \cup A_Y) \cap (B_X \cap B_Y)$ is a corner of X and Y. □

Lemma 5.39. *Let $X, Y \in \mathcal{A}$ be a crossing pair of separators and let (A_X, B_X) and (A_Y, B_Y), respectively, be witnesses that these are in $\mathcal{A}$. If $Z \in \mathcal{A}$ is nested with X, and each of the corner separations $(A_X, B_X) \vee (A_Y, B_Y)$ and $(A_X, B_X) \wedge (A_Y, B_Y)$ distinguishes some pair of profiles efficiently then Z is nested with one of the separators $(A_X \cup A_Y) \cap (B_X \cap B_Y)$ or $(A_X \cap A_Y) \cap (B_X \cup B_Y)$.*

Proof. Since Z and X are nested there is a component C^Z of $G - X$ such that $Z \subseteq C^Z \cup X$. Let us assume without loss of generality that $C^Z \subseteq A_X$, we will show that Z is nested with $(A_X \cup A_Y) \cap (B_X \cap B_Y)$.

Since $(A_X, B_X) \vee (A_Y, B_Y)$ efficiently distinguishes some regular profiles there is, by Lemma 5.26, a tight component of $(A_X \cup A_Y) \cap (B_X \cap B_Y)$ contained in $(B_X \cap B_Y)$. However, $Z \subseteq A_X$, so this component cannot meet Z. Hence, by Lemma 5.35, Z cannot cross the separator $(A_X \cup A_Y) \cap (B_X \cap B_Y)$. □

These now allow us to reuse Lemma 5.30 and 5.31 to prove conditions (TS2) and (TS3):

Lemma 5.40. *If two separators $X \in \mathcal{A}_{P,P'}$ and $Y \in \mathcal{A}_{Q,Q'}$ cross and we have $|P,P'| < |Q,Q'|$, then there is a corner $Y' \in \mathcal{A}_{Q,Q'}$ of X and Y which is nested with X.*

Proof. Let (A_X, B_X) be a separation witnessing that $X \in \mathcal{A}_{P,P'}$ and let (A_Y, B_Y) be a separation witnessing that $Y \in \mathcal{A}_{Q,Q'}$. By Lemma 5.30 there is a corner separation of (A_X, B_X) and (A_Y, B_Y) which also distinguishes Q and Q' efficiently. The separator Y' of this corner separation does not meet all tight components of $G - X$, so Y' is nested with X and thus is by Lemma 5.38 as desired. □

Lemma 5.41. *If two separators $X \in \mathcal{A}_{P,P'}$ and $Y \in \mathcal{A}_{Q,Q'}$ cross and we have $|P,P'| = |Q,Q'| = k$, then either there is a corner $Y' \in \mathcal{A}_{Q,Q'}$ of X and Y which has a strictly lower k-crossing number than Y, or there is a corner $X' \in \mathcal{A}_{P,P'}$ of X and Y which has strictly lower k-crossing number than X.*

Proof. By switching their roles if necessary we may assume that the k-crossing number of Y is at most the k-crossing number of X. Let (A_X, B_X) be a separation witnessing that $X \in \mathcal{A}_{P,P'}$ and let (A_Y, B_Y) be a separation witnessing that $Y \in \mathcal{A}_{Q,Q'}$. By Lemma 5.31 there is a corner separation of (A_X, B_X) and (A_Y, B_Y) which efficiently distinguishes P and P' and whose opposite corner separation efficiently distinguishes either P and P' or Q and Q'. Let us denote their separators as Z and Z' respectively.

By Lemma 5.38 and 5.39 and the fact that Z and Z' are nested with both X and Y we have that Z and Z' are corners of X and Y and that the sum of the k-crossing numbers of Z and Z' is strictly lower than the sum of the k-crossing numbers of X and Y.

Thus, if the k-crossing number of Z is strictly lower than the k-crossing number of X, we can take Z for X'. Otherwise we can infer that the k-crossing number of Z' is strictly lower than that of Y. Since we assumed in the beginning that the $|i|$-crossing number of Y is not greater than that of X. This proves the claim since we can then take Z' for X' or Y', depending. □

With this all the requirements of Lemma 22 are satisfied. Immediately we obtain the main result of this section:

Theorem 24. *Given a set of distinguishable robust regular profiles $\mathcal{P}$ of a graph G there exists a canonical nested set of separators efficiently distinguishing any pair of profiles in $\mathcal{P}$.*

□

As noted before, to be able to deduce Theorem 5.20 and [14, Remark 8.3] we restrict our set $\mathcal{P}$ to be a set of *principal* robust profiles. A k-profile P in G is *principal* if it contains for every set X of less than k vertices a separation of the form $(V(G) \setminus C, C \cup X)$ where C is a connected component of $G - X$. In particular, every principal profile is regular. Note that this notion of principal profiles is equivalent to the notion of 'profiles' in Carmesin's [10]; the term *principal profiles* comes from [14]. Observe that in locally finite graphs an inductive application of the profile property (P) shows that every profile is principal.

This restriction to principal profiles is necessary for Theorem 5.20, as Elm and Kurkofka [43, Corollary 3.4] have shown that there is a graph together with a set of (non-principal but robust and distinguishable) profiles, which do not permit the existence of a nested set of separations distinguishing all of them.

Nested sets of separations

If we restrict $\mathcal{P}$ to a set of principal profiles, the nested set of separators from Theorem 24 can be transformed into a nested set of *separations* which still distinguishes all the profiles in $\mathcal{P}$ if we give up on canonicity. This task is not entirely trivial.

The natural approach would be to take for each separator every one of the separations belonging to one of its tight components, i.e., the separation $(C \cup X, V \setminus C)$ for every tight component C of $G - X$. However, if the separators overlap the resulting set of separations might not be nested. The following recent result by Elm and Kurkofka states that we need to omit no more than one of the tight components for each separator to reclaim nestedness.

Theorem 5.42 ([43, Corollary 6.1]). *Suppose that $\mathcal{Y}$ is a principal collection of vertex sets in a connected graph G. Then there is a function $\mathcal{K}$ assigning to each $X \in \mathcal{Y}$ a subset $\mathcal{K}(X) \subseteq \mathcal{C}_X$ (the set $\mathcal{C}_X$ consists of the components of $G - X$ whose neighbourhoods are precisely equal to X) that misses at most one component from $\mathcal{C}_X$, such that the collection*

$$\{\, \{V \setminus K,\, X \cup K\} : X \in \mathcal{Y} \text{ and } K \in \mathcal{K}(X)\, \}$$

is nested.

Here, a principal collection of vertex sets is just a set $\mathcal{Y}$ of subsets of V such that, for every $X, Y \in \mathcal{Y}$, there is at most one component of $G - X$ which is met by Y. In particular, any nested set of separators is a principal collection of vertex sets.

Having for every separator all but one of these tight component separations is still enough to efficiently distinguish all the profiles in $\mathcal{P}$. However, as Theorem 5.42 does not give as a canonical choice for the function $\mathcal{K}$, we need to give up the canonicity at this point. However, this still allows us to prove the following theorem by Carmesin:

⇝ p. 97 **Theorem 5.20** ([10, Theorem 5.12]). *For any graph G, there is a nested set N of separations that distinguishes efficiently any two robust principal profiles (that are not restrictions of one another).*

Proof. If G is not connected, then every robust principal profile of G induces a robust principal profile on exactly one of the connected components of G. It is easy to see that we can then apply the theorem to all connected components from G independently and obtain our desired nested set of separations of G from those of the connected components together with separations of the form $(C, V \setminus C)$ for connected components C of G. Thus let us suppose that G is connected.

Let N be the nested set of separations obtained by applying Theorem 5.42 to the set $\mathcal{N}$ of separators obtained from Theorem 24. Given any two profiles $P, Q \in \mathcal{P}$ there is a separator X in $\mathcal{N}$ which efficiently distinguishes P and Q. By Lemma 5.26 there are two distinct tight components C and C' of $G - X$ such that both $(V \setminus C, C \cup X) \in P$ and $(C' \cup X, V \setminus C') \in P$ efficiently distinguish P and Q. However, at least one of these two separations is an element of N. □

For the reader's convenience, we also offer a direct proof of Theorem 5.20 which does not use Theorem 5.42. Instead, we perform an argument akin to one of the arguments used in the proof of Theorem 5.42 but in slightly simpler form, as the statement we need is a weaker one than Theorem 5.42.

Direct proof of Theorem 5.20. Let $\mathcal{N}$ be the nested set of separators obtained from Theorem 24 applied to the set of robust principal profiles. Pick an enumeration of $\mathcal{N}$ which is increasing in the size of the separators, i.e., an enumeration $\mathcal{N} = \{ X_\alpha : \alpha < \beta \}$ such that $|X_\alpha| \leqslant |X_\gamma|$ whenever $\alpha < \gamma$.

We will construct a transfinite ascending sequence of nested sets $(N_\gamma)_{\gamma \leqslant \beta}$, of separations. Each N_γ will contain only separations with separators in $\{ X_\alpha : \alpha < \gamma \}$, and every pair of profiles efficiently distinguished by such a separator X_α, $\alpha < \gamma$, will also be efficiently distinguished by some separation in N_γ.

For the successor steps of our construction suppose that we already constructed N_γ and consider X_γ. Since X_γ is nested with all X_α satisfying $\alpha < \gamma$ we know that X_γ induces a consistent orientation of N_γ since any separation $(A, B) \in N_\gamma$ satisfies either $X_\gamma \subseteq A$ or $X_\gamma \subseteq B$ but not both, as $|(A, B)| \leqslant |X_\gamma|$.

Consider the set $\mathcal{C}$ of tight components of $G - X_\gamma$ and let $\mathcal{D}$ be the set of the remaining, non-tight, components of $G - X_\gamma$.

Given any separation $(A, B) \in N_\gamma$ pointing away from X_γ (that is $X_\gamma \subseteq A$), the side B is contained in the union of one component $C_B \in \mathcal{C}$ together with some components in $\mathcal{D}$: Since X_γ is nested with $A \cap B$ there is a component in $G - X_\gamma$ containing $(A \cap B) \setminus X_\gamma$, thus, any other component C of $G - X_\gamma$ meeting B does not meet $A \cap B$ and must therefore satisfy $N(C) \subseteq A \cap B \cap X_\gamma$, i.e., this component is not tight.

Given a tight component $C \in \mathcal{C}$ let $\mathcal{D}_C \subseteq \mathcal{D}$ be the set of all components D in $\mathcal{D}$ with the property that there is some $(A, B) \in N_\gamma$ pointing away from X_γ such that D meets B and $C_B = C$. Informally, these sets $\mathcal{D}_C$ are the components which we will need to group together with their C when choosing our next separations. The $\mathcal{D}_C$ are pairwise disjoint: Indeed, given two separations (A, B) and (A', B') pointing away from X_γ, if $(B', A') \leqslant (A, B)$ then the set B' and B are disjoint, and if $(A, B) \leqslant (A', B')$, then $(A' \cap B') \setminus X_\gamma$ and $(A \cap B) \setminus X_\gamma$ cannot be contained in different tight components of $G - X_\gamma$.

Let N_{γ^+} consist of N_γ together with, for every tight component $C \in \mathcal{C}$ of $G - X_\gamma$, the separation $\Big(C \cup \bigcup \mathcal{D}_C \cup X_\gamma,\ V(G) \setminus (C \cup \bigcup \mathcal{D}_C)\Big)$. It is easy to see that this set is a nested set of separations. Moreover, any pair of profiles efficiently distinguished by X_γ is efficiently distinguished by one of these new separations.

For limit ordinals γ let $N_\gamma := \bigcup_{\alpha<\gamma} N_\alpha$, this set is nested since every pair in N_γ is already in some N_α.

Then $N := N_\beta$ is the desired nested set of separations. □

Canonical trees of tree-decompositions

To canonically and efficiently distinguish a robust set of principal profiles in a graph Carmesin, Hamann, and Miraftab [14] introduced more complex objects than nested sets of separations: trees of tree-decompositions. These consist of a rooted tree where every node is associated with a tree-decomposition. At the root this is a tree decomposition of G. At every remaining node there is a tree-decomposition of one of the torsos of the tree-decomposition at the parent node. Their main result is the following:

⇝ p. 97 **Theorem 5.21** ([14, Remark 8.3]). *Let G be a connected graph and $\mathcal{P}$ a distinguishable set of principal robust profiles in G. There exists a canonical tree of tree-decompositions with the following properties:*

(1) the tree of tree-decompositions distinguishes $\mathcal{P}$ efficiently;
(2) if $t \in V(T)$ has level k, then $(T_t, \mathcal{V}_t)$ contains only separations of order k;
(3) nodes t at all levels have $|V(T_t)|$ neighbours on the next level and the graphs assigned to them are all torsos of $(T_t, \mathcal{V}_t)$.

We can also construct such a tree of tree-decompositions from our nested set of separators. In order to do that, let us recall the most important definitions from [14].

In a rooted tree (T, r), the *level* of a vertex $t \in V(T)$ is $d(t, r) + 1$. A *tree of tree-decompositions* is a triple $((T, r), (G_t)_{t\in V(T)}, (T_t, \mathcal{V}_t)_{t\in V(T)})$ consisting of a rooted tree (T, r), a family $(G_t)_{t\in V(T)}$ of graphs and a family $(T_t, \mathcal{V}_t)_{t\in V(T)}$ of tree-decompositions of the G_t. The graphs G'_t assigned to the neighbours t' on the next level from a node $t \in V(T)$ shall be distinct torsos of the tree-decomposition $(T_t, \mathcal{V}_t)$. This tree of tree-decompositions is a tree of tree-decompositions *of* G, if $G_r = G$.

A separation (A, B) of G *induces* a separation (A', B') of G_t if $A \cap G_t = A'$ and $B \cap G_t = B'$. Given two profiles P, P', we say that a tree of tree-decompositions (*efficiently*) *distinguishes* P and P' if there is a separation (A, B) in G (efficiently) distinguishing them and a node $t \in V(T)$ such that the separation induced by (A, B) on G_t is one of the separation induced by the tree-decomposition $(T_t, \mathcal{V}_t)$ of G_t.

In order to deduce Theorem 5.21 from Theorem 24 it is useful to observe that our set of separators is nested in an even stronger sense: We say that two separators X and Y are *strongly nested* if there is a component C of $G-X$ such that $Y \subseteq C \cup N(C)$ and there is a component C' of $G-Y$ such that $X \subseteq C' \cup N(C')$. The separators from the nested set $\mathcal{N}$ from Theorem 24 are strongly nested:

Lemma 5.43. *If X and Y are a pair of nested separators each of which efficiently distinguishes some pair of robust principal profiles, then they are strongly nested.*

Proof. We show that there is a component C of $G-X$ such that $Y \subseteq C \cup N(C)$.

If $Y \subseteq X$ the statement is obvious, by picking as C a tight component of $G-X$. So we may assume that Y meets some component C of $G-X$ in a vertex $v \in Y \cap C$. By nestedness $Y \subseteq C \cup X$. Suppose for a contradiction that $Y \nsubseteq C \cup N(C)$, i.e., Y contains a vertex $w \in X \setminus N(C)$.

Since Y efficiently distinguishes two principal profiles there are two distinct tight components C_1, C_2 of $G-Y$, by Lemma 5.26. X meets at most one of C_1 and C_2 since it is nested with Y; without loss of generality we may assume $X \cap C_2 = \emptyset$. Since C_2 is a tight component of $G-Y$ there is a path P from v to w with all its interior vertices in C_2. On the other hand v lies in C and w outside of $C \cup N(C)$, so $N(C)$ separates v from w. But $N(C) \subseteq X$ does not meet P since $X \cap C_2 = \emptyset$. This is a contradiction. $\square$

Note that for a separator X to be strongly nested with itself is a non-trivial property: it is precisely the statement that there is a tight component of $G-X$. Thus, if we talk about a strongly nested set of separators, we mean that not only any pair of distinct separators from that set is strongly nested, we also require each of the separators from that set to be nested with itself.

Next we show that we can close our strongly nested set under taking subsets:

Lemma 5.44. *Let $\mathcal{N}$ be a strongly nested set of separators and let $\mathcal{N}'$ be the set of all subsets of elements of $\mathcal{N}$. Then $\mathcal{N}'$ is strongly nested as well.*

Proof. Let $X, Y \in \mathcal{N}$ and let $X' \subseteq X, Y' \subseteq Y$, possibly equal. Take C_X to be a component of $G-X$ for which $Y \subseteq C_X \cup N(C_X)$, then in particular $Y' \subseteq C_X \cup N(C_X)$. Since $X' \subseteq X$ there is some component $C_{X'} \supseteq C_X$ of $G-X'$, thus $Y' \subseteq C_{X'} \cup N(C'_X)$.

By symmetry, we also find a component $C_{Y'}$ such that $X' \subseteq C_{Y'} \cup N(C_{Y'})$ $\square$

So let $\mathcal{N}'$ be the strongly nested set of all subsets of separators from $\mathcal{N}$, the canonical nested set of separators from Theorem 24. As such, $\mathcal{N}'$ is canonical as well. The following lemma about separations with strongly nested separators will allow us to construct a tree of tree-decompositions from $\mathcal{N}'$ inductively, starting with the separators of lowest size.

Lemma 5.45. *If X, Y are distinct strongly nested separators and (A_X, B_X) and (A_Y, B_Y) are separations with separators X and Y respectively, such that $Y \subseteq B_X$,*

$X \subseteq B_Y$, then either (A_X, B_X) and (A_Y, B_Y) are nested, or there is a component C of $G - (X \cap Y)$ which meets neither X nor Y.

Proof. Suppose that $(A_Y, B_Y) \not\leq (B_X, A_X)$. Then either there is a vertex in A_Y which does not lie in B_X, or there is a vertex in A_X which does not lie in B_Y. Since $A_X \cap B_X = X \subseteq B_Y$ and $A_Y \cap B_Y = Y \subseteq B_X$ either of these cases implies that there is a vertex v in $(A_X \setminus B_X) \cap (A_Y \setminus B_Y)$. This vertex v needs to lie in some component C of $G - (X \cup Y)$. However, C cannot send an edge to $X \setminus Y$ since such an edge would contradict the fact that (A_Y, B_Y) is a separation. Similarly, C cannot be adjacent to any vertex of $Y \setminus X$. Thus C is in fact a component of $G - (X \cap Y)$ which meets either X nor Y. □

Now we are ready to deduce Theorem 5.21 from Theorem 24:

Proof of Theorem 5.21. Let $\mathcal{N}'$ be as above. We will build our tree of tree-decompositions inductively level-by-level, adding at stage k to every node t on level $k-1$ new neighbours on level k, one for every torso of the tree-decompositions $(T_t, \mathcal{V}_t)$. We do this in a way that ensures the following properties:

(a) If $d(r,t) = k$ then every separation in $(T_t, \mathcal{V}_t)$ has order $k+1$.

(b) Every separator in $\mathcal{N}'$ of size at least $k+2$ is contained in exactly one of the torsos of $(T_t, \mathcal{V}_t)$ whenever $d(t,r) \leqslant k$.

(c) If $d(t,r) = k$, every torso of $(T_t, \mathcal{V}_t)$ meets at most one component of $G - X$ for every $X \in \mathcal{N}'$ of size $\leqslant k$ with $X \subseteq V(G_t)$.

Our inductive construction goes as follows: For $k = 0$ we consider the set S_1 which consists of, for every separator X of size 1 in $\mathcal{N}'$ and every component C of $G - X$, the separation $(C \cup X, V(G) \setminus C)$, unless C is the only component of $G - X$.

Observe that S_1 is a nested set of separations: any two separations with the same separator are nested by construction and for separations with distinct separators X and Y the separators are disjoint, so $G - (X \cap Y) = G$ is connected and Lemma 5.45 gives that the separations are nested.

Moreover, every ω-chain $(A_1, B_1) < (A_2, B_2) < \ldots$ in S_1 has $\bigcap_{i \in \mathbb{N}} B_i = \emptyset$: We may assume without loss of generality that no two of these separations have the same separator since S_1 has no 3-chain of separations with the same separator. On the other hand a path from a vertex in $\bigcap_{i \in \mathbb{N}} B_i$ to A_1 (which has finite length) would need to meet all the infinitely many disjoint separators $A_i \cap B_i$.

Since S_1 contains no separation with a small orientation, it is, by construction, a regular tree set. Thus, by Lemma 5.27 it induces a canonical tree-decomposition $(T_r, \mathcal{V}_r)$ of $G_r = G$. We assign this tree-decomposition to the root of our tree of tree-decompositions and shall now verify properties (a) to (c).

Observe that this decomposition satisfies properties (a) and (c) as we only used separators of size 1 and every torso of $(T_t, \mathcal{V}_t)$ meets at most one component of

$G - X$ for every $X \in \mathcal{N}'$ of size $\leqslant 1$ with $X \subseteq V(G_t)$. Moreover, property (b) is also satisfied since every separator X in $\mathcal{N}'$ of size at least 2 is nested with each of the separators used in $(T_r, \mathcal{V}_r)$: Such a separator cannot be contained in two distinct torsos since then a separation with separator in $\mathcal{N}'$ would separate them. Conversely, there is a torso which contains X: Otherwise consider a torso V_t that contains as much of X as possible and another torso V_t' which contains a vertex in $X \setminus V_t$. Then one of the edges on the path between t and t' in T again corresponds to a separation which separates X. But this is not possible since the separators of these separations are in $\mathcal{N}'$ and thus nested with X.

For the k-th step of our construction, where $k \geqslant 1$, we attach at every node t on level $k-1$ of our so-far constructed tree of tree-decompositions, for every torso G' of $(T_t, \mathcal{V}_t)$ a new node t' (which then is at level k) with $G_{t'} := G'$. We the independently construct tree-decompositions for each of these torsos $G_{t'}$. For every torso we use all those separators from $\mathcal{N}'$ which are of size $k+1$ and lie inside that torso. Note that property (b) guarantees that every separator in $\mathcal{N}'$ of size $k+1$ is contained in exactly one of the newly added torsos.

Given one torso $G_{t'}$ of the tree-decomposition $(T_t, \mathcal{V}_t)$, we let S_{k+1} be the set of all separations (A, B) of $G_{t'}$ of order $k+1$ with separator in $\mathcal{N}'$ and the property that $A \setminus B$ is a component of $G - (A \cap B)$ but not the only one.

We claim that S_{k+1} is a nested set of separations. Indeed, if two separations from S_{k+1} with different separators X and Y were to cross, then by Lemma 5.45 there would be a component of $G_{t'} - (X \cap Y)$ avoiding X and Y. However, $X \cap Y$ has size less than k, lies in $\mathcal{N}'$ and $G_{t'}$ meets, by property (c), at most one component of $G - (X \cap Y)$. Hence, if we take vertices x and y in $G_{t'} - (X \cap Y)$, we find a path P between them in $G - (X \cap Y)$. But since $G_{t'}$ is obtained from G by repeatedly building a torso, $P \cap G_{t'}$ needs to contain a path between x and y in $G_{t'}$. In particular, this path does not meet $X \cap Y$ and thus $G_{t'} - (X \cap Y)$ has only one component, in particular every component of $G_{t'} - (X \cap Y)$ meets X and Y.

Now consider an ω-chain $(A_1, B_1) < (A_2, B_2) < \ldots$ in S_{k+1}. We may assume without loss of generality that no two of these separations have the same separator, as in the case $k = 0$. If $\bigcap_{i \in \mathbb{N}} B_i$ is non-empty then its neighbourhood $Z := N_{G_{t'}}(\bigcap_{i \in \mathbb{N}} B_i)$ needs to be properly contained in some $A_l \cap B_l$: Every vertex in Z needs to be contained in some $A_m \cap B_m$ and if such a vertex lies in $A_m \cap B_m$, then it also lies in $A_n \cap B_n$ for every $n \geqslant m$. In particular, if $|Z| \geqslant k+1$, there would be an m such that $A_m \cap B_m \subseteq Z$ and thus $A_n \cap B_n = A_m \cap B_m \forall n \geqslant m$ contradicting the assumption that no two of the (A_l, B_l) have the same separator. Hence, $|Z| \leqslant k$ and we can easily find an l so that $Z \subsetneq A_l \cap B_l$.

But then, again, $G_{t'}$ would meet two distinct components of $G - Z$: one meeting $\bigcap_{i \in \mathbb{N}} B_i$ and one meeting A_l. This, however, is not possible since $|Z| < |A_l \cap B_l|$ and $Z \in \mathcal{N}'$.

By construction S_{k+1} contains no separation with a small orientation, it is thus a regular tree set. So, by Lemma 5.27, the set S_{k+1} induces a canonical tree-decomposition $(T_{t'}, \mathcal{V}_{t'})$ of $G_{t'}$. In this way we construct all the tree-decompositions for nodes at level k. We need to verify properties (a) to (c). Property (a) is obvious. For property (b) we observe that every separator in $\mathcal{N}'$ of size at least $k+2$ which was contained in $G_{t'}$ was nested with every separator of a separation in S_{k+1} and is therefore contained in exactly one of the torsos of $(T_{t'}, \mathcal{V}_{t'})$, by the same argument as in the case $k = 0$.

For property (c) we note that for separators X of size $\leqslant k$ every torso of $(T_{t'}, \mathcal{V}_{t'})$ meets at most one component of $G - X$ as, by induction $G_{t'}$ itself only meets one component of $G - X$. For a separator X of size $k+1$ let H be a torso of $(T_{t'}, \mathcal{V}_{t'})$. Firstly, H meets at most one component of $G_{t'} - X$ since if $G_{t'} - X$ has more than one component then X is one of the separators of $(T_{t'}, \mathcal{V}_{t'})$ and therefore, as S_{k+1} includes every separation of the form $(C \cup X, G_{t'} \setminus X)$ for any component C of $G_{t'} - X$, there needs to be a component C of $G_{t'} - X$ such that H is contained in $C \cup X$.

Secondly, when building the torso $G_{t'}$ from G we never add edges between distinct components of $G - X$ since we only add edges inside of separators in $\mathcal{N}'$, which are nested with X. Hence, if H would meet two components of $G - X$ it would also meet two components of $G_{t'} - X$. Hence H meets at most one component of $G - X$. This gives property (c).

Verification of correctness. Let us now verify that the so constructed tree of tree-decompositions $((T, r), (G_t)_{t \in V(t)}, (T_t, \mathcal{V}_t)_{t \in V(T)})$ – which is canonical by construction – has the properties (1)–(3) from the assertion. The properties (2) and (3) are fulfilled by construction, so we only need to verify (1).

Let P, P' be two robust principal profiles from $\mathcal{P}$. By Theorem 24, $\mathcal{N}'$ contains some separator X which belongs to a separation efficiently distinguishing P and P', say $|X| = k$. By our inductive construction, there is a unique G_t at level k which contains X. As P and P' are principal profiles, there are two distinct components C, C' of $G - X$ such that $(V(G) \setminus C, C \cup X) \in P$, and $(V(G) \setminus C', C' \cup X) \in P'$. We claim that $C \cap V(G_t)$ is not empty.

Note that G_t is obtained from G by repeatedly taking some separation (A, B) of order $< k$ with $X \subseteq B$, deleting $A \setminus B$ and making $A \cap B$ complete. If we apply this operation for a single (A, B) which, say, turns some graph H with $V(H) \subseteq V(G)$ into H' then this preserves for H' the properties of H that (i) $H[C \cap V(H)]$ is connected and (ii) every vertex in X has, in H, a neighbour in $C \cap V(H)$. Thus every vertex in X has, in G_t, a neighbour in $C \cap V(G_t)$ proving that $C \cap V(G_t)$ is non-empty.

By a symmetrical argument not only C but also C' meets some component of $G_t - X$. Moreover, no two distinct components of $G - X$ can meet the same component of $G_t - X$: this would require an edge between these components, which would have to be added by the torso operation – but this operation only adds edges inside a separator $Y \in \mathcal{N}'$. And since Y is nested with X, that is Y meets only one component of $G - X$, this cannot add edges between different components of $G - X$.

Thus, there is exactly one component C_t of $G_t - X$ such that $C_t \subseteq C$ and this component is not the only one from $G_t - X$. So, by construction the separation $(C_t \cup X, G_t \setminus X)$, which efficiently distinguishes the induced profiles of P and P' onto G_t is induced by $(T_t, \mathcal{V}_t)$. $\square$

5.10. Distinguishing edge-blocks in finite and infinite graphs

The k-tangles and k-profiles in a graph that we considered in the previous section are a way of expressing its 'k-vertex-connected pieces', though these are very indirect ways of expressing 'pieces' of a graph via orientations of separations. A much more concrete structure are the *k-blocks* of a graph and, analogously to these, if we consider edge-connectivity instead of vertex-connectivity, there is a single concrete and undeniably natural notion of 'k-edge-connected pieces'. Let $k \in \mathbb{N} \cup \{\infty\}$ and let G be any connected graph, possibly infinite. We say that two vertices or ends are *$<k$-edge-inseparable* in G if they cannot be separated in G by fewer than k edges. This defines an equivalence relation on $\hat{V}(G) := V(G) \cup \Omega(G)$ where $\Omega(G)$ denotes the set of ends of G (which is empty if G is finite). Its equivalence classes are the 'k-edge-connected pieces' of G, its *k-edge-blocks*. A subset of $\hat{V}(G)$ is an *edge-block* if it is a k-edge-block for some k. Note that any two edge-blocks are either disjoint or one contains the other. The thin splinter lemma allows us to find a canonical tree-like decomposition of any connected graph, finite or infinite, into its k-edge-blocks — for all $k \in \mathbb{N} \cup \{\infty\}$ simultaneously. To state our result, we only need a few intuitive definitions.

A subset $X \subseteq \hat{V}(G)$ *lives in* a subgraph $C \subseteq G$ or vertex set $C \subseteq V(G)$ if all the vertices of X lie in C and all the rays of ends in X have tails in C or $G[C]$, respectively. If G is finite, saying that X lives in C simply means that $X \subseteq C$. An edge set $F \subseteq E(G)$ *distinguishes* two edge-blocks of G, not necessarily k-edge-blocks for the same k, if they live in distinct components of $G - F$. It distinguishes them *efficiently* if they are not distinguished by any edge set of smaller size. Note that if F distinguishes two edge-blocks efficiently, then F must be a *bond*, a cut with connected sides. A set B of bonds *distinguishes* some set of edge-blocks of G *efficiently* if every two disjoint edge-blocks in this set are distinguished efficiently by a bond in B. Two cuts F_1, F_2 of G are *nested* if F_1 has a side V_1 and F_2 has a side V_2 such that $V_1 \subseteq V_2$. Note that this is symmetric. The fundamental cuts of a spanning tree, for example, are (pairwise) nested. Our main result reads as follows:

Theorem 25. *Every connected graph G has a nested set of bonds that efficiently distinguishes all the edge-blocks of G.*

The nested sets $N = N(G)$ that we construct, one for every G, have two strong additional properties:

(i) They are canonical in that they are invariant under isomorphisms: if $\varphi\colon G \to G'$ is a graph-isomorphism, then $\varphi(N(G)) = N(\varphi(G))$.

(ii) For every $k \in \mathbb{N}$, the subset $N_k \subseteq N$ formed by the bonds of size less than k is equal to the set of fundamental cuts of a tree-cut decomposition of G that decomposes G into its k-edge-blocks.

Tree-cut decompositions are decompositions of graphs similar to tree-decompositions but based on edge-cuts rather than vertex-separators. They were introduced by Wollan [69], and they are more general than the 'tree-partitions' introduced by Seese [65] and by Halin [55]; see Section 5.10.2.

The second additional property above is best possible in the sense that N_k cannot be replaced with N: there exists a graph G (see Example 5.55) that has no nested set of cuts which, on the one hand, distinguishes all the edge-blocks of G efficiently, and on the other hand, is the set of fundamental cuts of some tree-cut decomposition. (This is because the 'tree-structure' defined by a nested set of cuts may have limit points, and hence not be representable by a graph-theoretical tree.)

It turns out that the nested sets of bonds which make Theorem 25 true can be characterised in terms of generating bonds (for the definition of *generate* see Section 5.10.3):

Theorem 26. *Let G be any connected graph and let M be any nested set of bonds of G. Then the following assertions are equivalent:* ⇝ p. 126

(i) M efficiently distinguishes all the edge-blocks of G;

(ii) For every $k \in \mathbb{N}$, the $\leqslant k$-sized bonds in M generate all the k-sized cuts of G.

Nested sets of bonds which are canonical and satisfy assertion (ii) of Theorem 26 have been constructed by Dicks and Dunwoody using their algebraic theory of graph symmetries. This is one of the main results of their monograph [17, II 2.20f]. Since the implication (ii)→(i) of Theorem 26 is straightforward, Theorem 25 can be deduced from their theory, but it is not stated in [17] explicitly. Our Theorem 26 itself, in particular its highly non-trivial forward implication (i)→(ii), does not follow from material in [17]. Since our proofs are purely combinatorial, we can combine Theorem 25 and the forward implication (i)→(ii) of Theorem 26 to obtain a purely combinatorial proof of the main result of Dicks and Dunwoody.

After setting up just a bit more terminology and tools, we will prove Theorem 25 using Lemma 22 in Section 5.10.1. In Section 5.10.2 we relate each N_k to a tree-cut decomposition. In Section 5.10.3 we prove Theorem 26.

Additional terminology and tools

Throughout this chapter, $G = (V, E)$ denotes a fixed connected graph, finite or infinite. When we say ends we mean the usual vertex-ends, not edge-ends. If a subset $X \subseteq \hat{V}(G)$, usually an edge-block, lives in a subgraph $C \subseteq G$ or vertex set $C \subseteq V(G)$, we denote this by $X \sqsubseteq C$ for short. Recall that $X \sqsubseteq C$ defaults to $X \subseteq C$ if G is finite.

The *order* of a cut is its size. A *cut-separation* of a graph G is a bipartition $\{A, B\}$ of the vertex set of G, and it *induces* the cut $E(A, B)$. Then the order of the cut $E(A, B)$ is also the *order* of $\{A, B\}$. Recall that in a connected graph, every cut is induced by a unique cut-separation in this way, to which it *corresponds*. A *bond-separation* of G is a cut-separation that induces a bond of G, a cut with connected sides. We say that a cut-separation *distinguishes* two edge-blocks (*efficiently*) if its corresponding cut does, and we call two cut-separations *nested* if their corresponding cuts are nested. Thus, two cut-separations $\{A, B\}$ and $\{C, D\}$ are nested if one of the four inclusions $A \subseteq C$, $A \subseteq D$, $B \subseteq C$ or $B \subseteq D$ holds.

We use the following lemma about bonds which is well-known [19, Exercise 8.12]; we provide a proof for the reader's convenience.

Lemma 5.46. *Every edge of a graph lies in only finitely many bonds of size k of that graph, for any $k \in \mathbb{N}$.*

Proof. Let e be any edge of a graph G, and suppose for a contradiction that e lies in infinitely many distinct bonds $B_0, B_1, \dots$ of size k, say. Let F be an inclusionwise maximal set of edges of G such that F is included in B_n for infinitely many n (all n, without loss of generality). Then $|F| < k$ because the bonds are distinct, and any bond $B_n \supsetneq F$ gives rise to a path P in $G - F$ that links the endvertices of e. Now all the infinitely many bonds B_n must contain an edge of the finite path P. But by the choice of F, each edge of P lies in only finitely many B_n, a contradiction. $\square$

Corollary 5.47. *Let G be any connected graph, $k \in \mathbb{N}$, and let $F_0, F_1, \dots$ be infinitely many distinct bonds of G of size at most k such that each bond F_n has a side A_n with $A_n \subsetneq A_m$ for all $n < m$. Then $\bigcup_{n \in \mathbb{N}} A_n = V$.*

Proof. If $\bigcup_n A_n$ is a proper subset of V, then any A_0–$(V \setminus \bigcup_n A_n)$-path in G admits an edge that lies in infinitely many F_n, contradicting Lemma 5.46. $\square$

5.10.1. Proof of Theorem 25

Let G be any connected graph, possibly infinite, and consider the set $\mathcal{A}$ of all bond-separations of G with the relation $\sim$ being the usual nestedness-realation on bipartitions. For the family $(\mathcal{A}_i : i \in I)$ we will consider I to be the collection of all the unordered pairs formed by two disjoint edge-blocks of G, and each $\mathcal{A}_i$ will consist of all the bond-separations of G that efficiently distinguish the two edge-blocks forming the pair i. Our choice for $|i|$ will be the unique natural number that is the order of all the bond-separations in $\mathcal{A}_i$. Note that each of the two edge-blocks forming the pair i will be a k-edge-block for some $k > |i|$.

Our aim is to employ Lemma 22 to deduce Theorem 25. In order to do that, we first have to verify that $(\mathcal{A}_i : i \in I)$ thinly splinters. To this end, we verify all the three properties (TS1)–(TS3) below. The following lemma clearly implies condition (TS1):

Lemma 5.48. *Every finite-order bond-separation of a graph G is crossed by only finitely many bond-separations of G of order at most k, for any given $k \in \mathbb{N}$.*

Proof. Our proof starts with an observation. If two bond-separations $\{A, B\}$ and $\{A', B'\}$ cross, then A' contains a vertex from A and a vertex from B. Let $v \in A' \cap A$ and $w \in A' \cap B$. Since $G[A']$ is connected, there exists a path from v to w in $G[A']$. This path, and thus $G[A']$, must contain an edge from A to B. Similarly, $G[B']$ must contain an edge from A to B.

Now suppose for a contradiction that there are infinitely many bond-separations of order at most a given $k \in \mathbb{N}$, which all cross some finite-order bond-separation $\{A, B\}$. Without loss of generality, all the crossing bond-separations have order k. Using our observation, the pigeon-hole principle and the finite order of $\{A, B\}$, we find two edges $e, f \in E(A, B)$ and infinitely many bond-separations $\{A_0, B_0\}, \{A_1, B_1\}, \ldots$ that all cross $\{A, B\}$ so that $e \in G[A_n]$ and $f \in G[B_n]$ for all $n \in \mathbb{N}$. Let P be a path in G that links an endvertex v of e to an endvertex w of f. Now v is contained in all the A_n and w is contained in all the B_n, thus for every $\{A_n, B_n\}$ there exists an edge of P with one end in A_n and the other in B_n. However, every $\{A_n, B_n\}$ corresponds to a bond of size k of G and, again by the pigeon-hole principle, infinitely many of these bonds must contain the same edge of P. This contradicts Lemma 5.46. □

Next, to show the second property, we use the following lemma special case of the fish lemma, Lemma 2.1.

Lemma 5.49. *If two cut-separations $\{A_1, B_1\}$ and $\{A_2, B_2\}$ cross, and a third cut-separation $\{X, Y\}$ is nested with both $\{A_1, B_1\}$ and $\{A_2, B_2\}$, then $\{X, Y\}$ is nested with $\{A_1 \cap A_2, B_1 \cup B_2\}$ (provided that this is a cut-separation).*

Using this, we now show condition (TS2):

Lemma 5.50. *If $\{A, B\} \in \mathcal{A}_i$ and $\{C, D\} \in \mathcal{A}_j$ cross with $|i| < |j|$, then $\mathcal{A}_j$ contains some corner of $\{A, B\}$ and $\{C, D\}$ that is nested with $\{A, B\}$.*

Proof. Let us denote the two edge-blocks in j as β and β' so that $\beta \sqsubseteq C$ and $\beta' \sqsubseteq D$. Since the order of $\{A, B\}$ is less than $|j|$, we may assume without loss of generality that $\beta, \beta' \sqsubseteq A$. We claim that either $\{A \cap C, B \cup D\}$ or $\{A \cap D, B \cup C\}$ is the desired corner in $\mathcal{A}_j$, and we refer to them as *corner candidates*. Both are cut-separations that distinguish β and β', and both are nested with $\{A, B\}$. Furthermore, by Lemma 5.49, every cut-separation that is nested with both $\{A, B\}$ and $\{C, D\}$ is also nested with both corner candidates. It remains to show that at least one of the two corner candidates has order at most $|j|$, because then it lies in $\mathcal{A}_j$ as desired.

Let us assume for a contradiction that both corner candidates have order greater than $|j|$. Then the two inequalities

$$\begin{aligned} & |E(A \cap C, B \cup D)| + |E(B \cap D, A \cup C)| \leqslant |E(A, B)| + |E(C, D)| \\ \text{and} \quad & |E(A \cap D, B \cup C)| + |E(B \cap C, A \cup D)| \leqslant |E(A, B)| + |E(C, D)| \end{aligned}$$

imply

$$|E(B \cap D, A \cup C)| < |i| \quad \text{and} \quad |E(B \cap C, A \cup D)| < |i|\,.$$

Recall that the edge-blocks forming the pair i are k-edge-blocks for some values k greater than $|i|$. One of the edge-blocks of the pair i lives in B, and due to the latter two inequalities, this edge-block must live either in $B \cap D$ or in $B \cap C$. But then either $\{ B \cap D, A \cup C \}$ or $\{ B \cap C, A \cup D \}$ is a cut-separation of order less than $|i|$ that distinguishes the two edge-blocks forming the pair i, contradicting the fact that an order of at least $|i|$ is required for that. □

Finally, to show the third property, we need the following lemma:

Lemma 5.51. *Let $\{ A_1, B_1 \}$ and $\{ A_2, B_2 \}$ be crossing cut-separations such that both $\{ A_1 \cap A_2, B_1 \cup B_2 \}$ and $\{ A_1 \cup A_2, B_1 \cap B_2 \}$ are cut-separations as well. Then every cut-separation that crosses both $\{ A_1 \cap A_2, B_1 \cup B_2 \}$ and $\{ A_1 \cup A_2, B_1 \cap B_2 \}$ must also cross both $\{ A_1, B_1 \}$ and $\{ A_2, B_2 \}$.*

Proof. Consider any cut-separation $\{ X, Y \}$ that crosses both $\{ A_1 \cap A_2, B_1 \cup B_2 \}$ and $\{ A_1 \cup A_2, B_1 \cap B_2 \}$. Since $\{ X, Y \}$ crosses $\{ A_1 \cap A_2, B_1 \cup B_2 \}$, both X and Y contain a vertex from $A_1 \cap A_2$. Since $\{ X, Y \}$ crosses $\{ A_1 \cup A_2, B_1 \cap B_2 \}$, both X and Y contain a vertex from $B_1 \cap B_2$. Hence $\{ X, Y \}$ crosses both $\{ A_1, B_1 \}$ and $\{ A_2, B_2 \}$. □

Let us now show condition (TS3):

Lemma 5.52. *If $\{ A, B \} \in \mathcal{A}_i$ and $\{ C, D \} \in \mathcal{A}_j$ cross with $|i| = |j| = k \in \mathbb{N}$, then either $\mathcal{A}_i$ contains a corner of $\{ A, B \}$ and $\{ C, D \}$ with strictly lower k-crossing number than $\{ A, B \}$, or else $\mathcal{A}_j$ contains a corner of $\{ A, B \}$ and $\{ C, D \}$ with strictly lower k-crossing number than $\{ C, D \}$.*

Proof. Let us assume without loss of generality that the k-crossing number of $\{ A, B \}$ is less than or equal to the k-crossing number of $\{ C, D \}$, and let us denote the edge-blocks in j as β and β' so that $\beta \sqsubseteq C$ and $\beta' \sqsubseteq D$. We consider two cases.

In the first case, $\{ A, B \}$ distinguishes the two edge-blocks β and β'. Hence $\beta \sqsubseteq A \cap C$ and $\beta' \sqsubseteq B \cap D$, say. Then both $\{ A \cap C, B \cup D \}$ and $\{ B \cap D, A \cup C \}$ distinguish the two edge-blocks β and β' that form the pair j, and so they have order at least $|j| = k$. Furthermore, we have submodularity:

$$|E(A \cap C, B \cup D)| + |E(B \cap D, A \cup C)| \leqslant |E(A, B)| + |E(C, D)| = 2k, \tag{5.1}$$

so both $\{ A \cap C, B \cup D \}$ and $\{ B \cap D, A \cup C \}$ must have order exactly k. In particular, both are contained in $\mathcal{A}_j$, and they are corners of $\{ A, B \}$ and $\{ C, D \}$ by Lemma 5.49. Next, we assert that the k-crossing numbers of $\{ A \cap C, B \cup D \}$ and $\{ B \cap D, A \cup C \}$ in sum are less than the sum of the k-crossing numbers of $\{ A, B \}$ and $\{ C, D \}$. Indeed, all the k-crossing numbers involved are finite by condition (TS1), and the two cut-separations $\{ A, B \}$ and $\{ C, D \}$ cross which allows us to deduce the desired inequality between the sums by Lemma 5.49 and 5.51, as follows:

- by Lemma 5.49, every $\{ X, Y \} \in \mathcal{A}$ of order k that crosses at least one of $\{ A \cap C, B \cup D \}$ and $\{ B \cap D, A \cup C \}$ must cross at least one of $\{ A, B \}$ and $\{ C, D \}$; and
- by Lemma 5.51, every $\{ X, Y \} \in \mathcal{A}$ of order k that crosses both $\{ A \cap C, B \cup D \}$ and $\{ B \cap D, A \cup C \}$ must cross both $\{ A, B \}$ and $\{ C, D \}$.

But then the strict inequality between the sums, plus our initial assumption that the k-crossing number of $\{ A, B \}$ is less than or equal to that of $\{ C, D \}$, implies that one of $\{ A \cap C, B \cup D \}$ and $\{ B \cap D, A \cup C \}$ must have a k-crossing number less than the one of $\{ C, D \}$, as desired.

In the second case, $\{ A, B \}$ does not distinguish the two edge-blocks β and β'. Recall that all the edge-blocks in the two pairs i and j are ℓ-edge-blocks for some values $\ell > k$. Hence $\beta \cup \beta' \sqsubseteq A$, say. Let us denote by β'' the edge-block in i that lives in B. Then either $\beta'' \sqsubseteq B \cap C$ or $\beta'' \sqsubseteq B \cap D$, say $\beta'' \sqsubseteq B \cap D$. In total:

$$\beta \sqsubseteq A \cap C, \ \ \beta' \sqsubseteq A \cap D \text{ and } \beta'' \sqsubseteq B \cap D.$$

Therefore, $\{ A \cap C, B \cup D \}$ distinguishes the two edge-blocks β and β' forming the pair j which imposes an order of at least k, and $\{ B \cap D, A \cup C \}$ distinguishes the two edge-blocks forming the pair i which imposes an order of at least k as well. Combining these lower bounds with (5.1) we deduce that both $\{ A \cap C, B \cup D \}$ and $\{ B \cap D, A \cup C \}$ have order exactly k. In particular, they are contained in $\mathcal{A}_j$ and $\mathcal{A}_i$ respectively, and they are corners of $\{ A, B \}$ and $\{ C, D \}$ by Lemma 5.49. Repeating the final argument of the first case, we deduce from the strict inequality between the sums of the k-crossing numbers that either $\{ A \cap C, B \cup D \} \in \mathcal{A}_j$ has strictly lower k-crossing number than $\{ C, D \}$, or else $\{ B \cap D, A \cup C \} \in \mathcal{A}_i$ has strictly lower k-crossing number than $\{ A, B \}$, completing the proof. □

With all three properties shown, the proof of our main result is all but complete:

Proof of Theorem 25. Let G be any connected graph. By Lemma 5.48, Lemma 5.50 and Lemma 5.52 we may apply Lemma 22 to the family $(\mathcal{A}_i : i \in I)$ defined at the beginning of the section. This results in the desired nested set $N(G) \subseteq \mathcal{A}$. To see that it is canonical, note that any isomorphism $\varphi \colon G \to G'$ induces an isomorphism between $(\mathcal{A}, \sim)$ and $(\mathcal{A}', \sim')$, where the latter is defined like the former but with regard to G'. Thus, by the 'moreover' part of Lemma 22, we indeed obtain that $\varphi(N(G)) = N(\varphi(G))$. □

5.10.2. From nested sets of bonds to tree-cut-decompositions

Recall that, given a connected graph G, we denote by $N = N(G)$ the canonical set of nested bonds from Theorem 25 that efficiently distinguishes all the edge-blocks of G. Furthermore, recall that the subset $N_k \subseteq N$ is formed by the bonds in N of order less than k. In this section, we show property (ii) from above, that is:

- For every $k \in \mathbb{N}$, the subset $N_k \subseteq N$ is equal to the set of fundamental cuts of a tree-cut decomposition of G that decomposes G into its k-edge-blocks.

To this end, we first introduce the notion of a tree-cut decomposition. Recall that a *near-partition* of a set M is a family of pairwise disjoint subsets $M_\xi \subseteq M$, possibly empty, such that $\bigcup_\xi M_\xi = M$.

Let G be a graph, T a tree, and let $\mathcal{X} = (X_t)_{t \in T}$ be a family of vertex sets $X_t \subseteq V(G)$ indexed by the nodes t of T. The pair $(T, \mathcal{X})$ is called a *tree-cut-decomposition* of G if $\mathcal{X}$ is a near-partition of $V(G)$. The vertex sets X_t are the *parts* or *bags* of the tree-cut-decomposition $(T, \mathcal{X})$. When we say that $(T, \mathcal{X})$ *decomposes* G into its k-edge-blocks for a given k, we mean that the non-empty parts of $(T, \mathcal{X})$ are the sets of vertices of the k-edge-blocks of G. For our purposes, we require the nodes with non-empty parts to be *dense* in T in that every edge of T lies on a path in T that links up two nodes with non-empty parts.

If $(T, \mathcal{X})$ is a tree-cut-decomposition, then every edge $t_1 t_2$ of its *decomposition tree* T induces a cut $E(\bigcup_{t \in T_1} X_t \,,\, \bigcup_{t \in T_2} X_t\,)$ of G where T_1 and T_2 are the two components of $T - t_1 t_2$ with $t_1 \in T_1$ and $t_2 \in T_2$. Here, the nodes with non-empty parts densely lying in T ensures that both unions are non-empty, which is required of the sides of a cut. We call these induced cuts the *fundamental cuts* of the tree-cut-decomposition $(T, \mathcal{X})$. Note that, unlike the fundamental cuts of a spanning tree, the fundamental cuts of a tree-cut-decomposition need not be bonds.

It is important that parts of a tree-cut-decomposition are allowed to be empty, as the following example demonstrates.

Example 5.53. Let the graph G arise from the disjoint union of three copies G_1, G_2 and G_3 of K^4 by selecting one vertex $v_i \in G_i$ for all $i \in [3]$ and adding all edges $v_i v_j$ (for $i \neq j \in [3]$). Then the 3-edge-blocks of G are the three vertex sets $V(G_1)$, $V(G_2)$ and $V(G_3)$. Since $N(G)$ is canonical, we have $N_3(G) = \{\, F_1, F_2, F_3 \,\}$ where $F_i := \{\, v_i v_j : j \neq i \,\}$. However, we cannot find a tree-cut-decomposition $(T, \mathcal{X})$ of G so that, on the one hand, T is a tree on three nodes t_1, t_2, t_3 and $X_{t_i} = V(G_i)$ for all $i \in [3]$, and on the other hand, the fundamental cuts of $(T, \mathcal{X})$ are precisely the bonds in $N_3(G)$: the decomposition tree T would then be a path of length two, and hence would induce two fundamental cuts, but $N_3(G)$ consists of three bonds.

To relate N_k to a tree-cut decomposition, we will use Theorem 2.2 by Gollin and Kneip. For this, we first note that cut-separations of a graph are an instance of set-separations or, more specifically, they form the separation system of bipartitions of $V(G)$. Similarly to the correspondence of tree-decompositions of graphs and nested sets of vertex separations (cf. Section 2.6.1), a tree-cut decomposition $(T, \mathcal{X})$ makes T into an order-respecting S-tree where S is the set of cut-separations which correspond to its fundamental cuts. We will now consider at the converse.

N_k is a set of fundamental cuts

The following theorem clearly implies that N_k is the set of fundamental cuts of a tree-cut decomposition of G that decomposes G into its k-edge-blocks:

Theorem 5.54. *Let G be any connected graph and $k \in \mathbb{N}$. Every nested set of bonds of G of order less than k is the set of fundamental cuts of some tree-cut-decomposition of G.*

Proof. Let G be any connected graph, $k \in \mathbb{N}$, and let B be any nested set of bonds of G of order less than k. We write S for the set of bond-separations which correspond to the bonds in B.

First, we wish to use Theorem 2.2 to find an S-tree (T, α) so that $\alpha \colon \vec{E}(T) \to \vec{S}$ is an isomorphism. For this, it suffices to show that B cannot contain pairwise distinct bonds $F_0, F_1, \dots, F_\omega$ such that each bond F_α has a side A_α with $A_\alpha \subsetneq A_\beta$ for all $\alpha < \beta \leqslant \omega$. This is immediate from Corollary 5.47.

Second, we wish to find a tree-cut-decomposition $(T, \mathcal{X})$ whose fundamental cuts are precisely equal to the bonds in B. We define the parts X_t of $(T, \mathcal{X})$ by letting

$$X_t := \bigcap \{\, D : (C, D) = \alpha(x, t) \text{ where } xt \in E(T) \,\}.$$

Then clearly the parts X_t are pairwise disjoint. To see that $\bigcup_t X_t$ includes the whole vertex set of G, consider any vertex $v \in V(G)$. We orient each edge $t_1 t_2 \in T$ towards the t_i with $v \in D$ for $(C, D) = \alpha(t_{3-i}, t_i)$. By Corollary 5.47 we may let t be the last node of a maximal directed path in T; then all the edges of T at t are oriented towards t, and $v \in X_t$ follows. Therefore, $\mathcal{X}$ is a near-partition of $V(G)$. It is straightforward to see that B is the set of fundamental cuts of $(T, \mathcal{X})$. □

N is not a set of fundamental cuts

Finally, we show that there exists a graph G that has no nested set of cuts which, on the one hand, distinguishes all the edge-blocks of G efficiently, and on the other hand, is the set of fundamental cuts of some tree-cut decomposition.

Example 5.55. This example is a variation of Example 5.33. Consider the locally finite graph displayed in Fig. 5.4. This graph G is constructed as follows. For every $n \in \mathbb{N}_{\geqslant 1}$ we pick a copy of $K^{2^{n+2}}$ together with $n+2$ additional vertices $w_1^n, \dots, w_{n+2}^n$. Then we select 2^n vertices of the $K^{2^{n+2}}$ and call them $u_1^n, \dots, u_{2^n}^n$. Furthermore, we select 2^{n+1} vertices of the $K^{2^{n+2}}$, other than the previously chosen u_i^n, and call them $v_1^n, \dots, v_{2^{n+1}}^n$. Now we add all the red edges $v_i^n u_i^{n+1}$, all the blue edges $w_i^n w_j^{n+1}$, and if $n \geqslant 2$ we also add the black edge $u_1^n w_1^n$. Finally, we disjointly add one copy of K^{10} and join one vertex v_1^0 of this K^{10} to u_1^1 and u_2^1; and we select another vertex $w_1^0 \in K^{10}$ distinct from v_1^0 and add all edges $w_1^0 w_i^1$. This completes the construction.

Now the vertex sets of the chosen $K^{2^{n+2}}$ are $(2^{n+2} - 1)$-edge-blocks B_n. The cut-separation $F_n := \{\bigcup_{k=1}^n B_n \,,\, V \setminus \bigcup_{k=1}^n B_n\}$ is the only one that efficiently

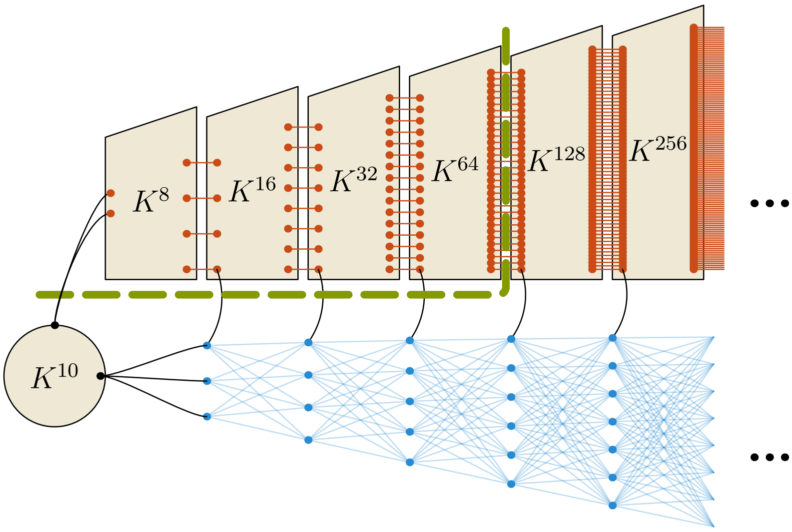

Figure 5.4.: The only cut that efficiently distinguishes the two edge-blocks defined by K^{64} and by K^{128} is drawn in green.

distinguishes B_n from B_{n+1}. Additionally, the vertex set of the K^{10} is a 9-edge-block B_0. The only cut-separation efficiently distinguishing B_0 and B_1 is $F_0 := \{B_0, V \setminus B_0\}$. Therefore, $N(G)$ must contain all the cuts corresponding to the cut-separations F_n $(n \in \mathbb{N})$. But the cut-separations F_n define an $(\omega + 1)$-chain

$$(B_1, V \setminus B_1) < (B_1 \cup B_2, V \setminus (B_1 \cup B_2)) < \cdots < (V \setminus B_0, B_0),$$

so $N(G)$ cannot be equal to the set of fundamental cuts of a tree cut-decomposition of G by Theorem 2.2.

5.10.3. Generating all bonds

A set S of cut-separations *generates* a cut $\{X, Y\}$ if and only if both (X, Y) and (Y, X) can be obtained from finitely many oriented cut-separations in $\vec{S}$ by taking suprema and infima in the universe of all cut-separations, i.e., the universe of bipartitions of $V(G)$. If S generates $\{X, Y\}$, then the cuts corresponding to the cut-separations in S *generate* the cut corresponding to $\{X, Y\}$. In this section we prove the following result:

Theorem 26. *Let G be any connected graph and let M be any nested set of bonds of G. Then the following assertions are equivalent:*

(i) M efficiently distinguishes all the edge-blocks of G;

(ii) For every $k \in \mathbb{N}$, the $\leqslant k$-sized bonds in M generate all the k-sized cuts of G.

For the proof, we use a generalized version of the star-comb lemma [19, Lemma 8.2.2]. A *comb* in a given graph G, for our generalization, shall mean one of the following two substructures of G:

1. The union of a ray R (the comb's *spine*) with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R. The last vertices of those paths are the *teeth* of this comb.

2. The union of a ray R (the comb's *spine*) with infinitely many disjoint pairwise inequivalent rays $R_0, R_1, \ldots$ that have precisely their first vertex on R. The ends to which the rays $R_0, R_1, \ldots$ belong are the *teeth* of this comb.

Given a set $U \subseteq V(G) \cup \Omega(G)$, a *comb attached to* U is a comb with all its teeth in U. A *star attached to* U is either a subdivided infinite star with all its leaves in U, or a union of infinitely many rays that meet precisely in their first vertex and belong to distinct ends in U.

Lemma 5.56 (Generalized star-comb lemma)**.** *Let $U \subseteq V(G) \cup \Omega(G)$ be an infinite set for a connected graph G. Then G contains either a comb attached to U or a star attached to U.*

Proof. If U contains infinitely many vertices of G, then we are done by the standard star-comb lemma [19, Lemma 8.2.2]. Hence, we may assume that U consists of ends and, say, is countable. Inductively, we choose for each end $\omega \in U$ a ray $R_\omega \in \omega$ so that R_ω is disjoint from all previously chosen rays, ensuring that all chosen rays are pairwise disjoint, and we let U' consist of the first vertices of these rays. Then we consider an inclusionwise minimal tree $T \subseteq G$ that extends all the rays R_ω with $\omega \in U$. Let $T' \subseteq T$ be the inclusionwise minimal subtree that contains U'. Then, by the standard star-comb lemma, T' contains either a star or a comb attached to U', and either extends to a star or comb attached to U. $\square$

For more on stars and combs, see the series [5–8].

Proof of Theorem 26. (ii)→(i). Let M be any nested set of bonds of G such that, for every $k \in \mathbb{N}$, the $\leqslant k$-sized bonds in M generate all the k-sized cuts of G, and suppose for a contradiction that there are two edge-blocks β_1, β_2 which are not efficiently distinguished by any bond in M. Let $\{X, Y\}$ be some bond-separation which efficiently distinguishes β_1 and β_2, and let k be its order. Let $\{\{A_\ell, B_\ell\} : \ell < n\}$ be a finite set of $\leqslant k$-sized bonds in M which generate $\{X, Y\}$. Since M does not efficiently distinguish β_1 from β_2, for every $\ell < n$ we either have that both β_1 and β_2 live in A_ℓ, or that both of them live in B_ℓ. However, this implies that either both β_1 and β_2 live in X, or that both of them live in Y, contradicting the fact that $\{X, Y\}$ distinguishes β_1 and β_2.

(i)→(ii). We assume (i). It suffices to prove (ii) for finite bonds. Let $B = E(V_1, V_2)$ be any bond of G of size k, say. By Theorem 5.54, the set formed by the $\leqslant k$-sized

bonds in M is the set of fundamental cuts of a tree-cut decomposition $(T, \mathcal{X})$ of G. Write (T, α) for the S-tree that arises from $(T, \mathcal{X})$.

Since B is finite, only finitely many parts of $(T, \mathcal{X})$ contain endvertices of edges in B. We let H be the minimal subtree of T which contains all the nodes corresponding to these parts. Note that H is finite. Then we let H' be the subtree of T which is induced by the nodes of H and all their neighbours in T. The subtree H' might be infinite, but it is rayless. Let $\mathcal{H}$ be the tree-cut decomposition of G which corresponds to the restricted S-tree $(H', \alpha \restriction \vec{E}(H'))$.

We claim that every two edge-blocks of G that are distinguished by B are also distinguished by some fundamental cut of $\mathcal{H}$. For this, let $\beta_1 \sqsubseteq V_1$ and $\beta_2 \sqsubseteq V_2$ be any two edge-blocks of G that are distinguished by B. Then β_1 and β_2 are also distinguished by a $\leqslant k$-sized bond in M, and hence some fundamental cut of $(T, \mathcal{X})$ distinguishes β_1 and β_2 as well. Let st be an edge of T whose induced fundamental cut distinguishes β_1 and β_2, chosen at minimal distance to H' in T. Then β_1 lives in C and β_2 lives in D for $(C, D) = \alpha(s, t)$, say. We claim that st is also an edge of H', and assume for a contradiction that it is not. Then s, say, is not a vertex of H' and t lies on the s–H'-path in T. Since $\{C, D\}$ is an element of M, it is a bond and in particular $G[C]$ is connected. Moreover, C avoids the endvertices of the edges in B, because t separates s from H. Therefore, C is included in one of the two sides of B, say in V_1, so β_1 lives in V_1. The node t, however, cannot lie in H because this would imply $s \in H'$, so t has a neighbour u in T which separates t (and s) from H. Let $(C', D') := \alpha(t, u)$. Since s and u are distinct neighbours of t, we have $(C, D) \leqslant (C', D')$. As argued for (C, D), we find that C' must be included in one of the two sides of B, and this side must be V_1 since C is included in both V_1 and C'. By the choice of st at minimal distance to H', the edge-block β_2 must live in C' (or we could replace st with tu, contradicting the choice of st). But then both β_1 and β_2 live in V_1, the desired contradiction.

We replace $(T, \mathcal{X})$ with $\mathcal{H}$. Then:

Every two edge-blocks of G that are distinguished by B are also distinguished by some fundamental cut of $(T, \mathcal{X})$. $(*)$

Given a node $t \in T$, we denote by $\hat{X}_t$ the subset of $\hat{V}(G)$ which is the union of all the $(k+1)$-edge-blocks of G that live in D for all cut-separations $(C, D) = \alpha(s, t)$ with $(s, t) \in \vec{E}(T)$. Then $\hat{X}_t \cap V(G) = X_t$ and we call $\hat{X}_t$ the *extended part* of t. Note that extended parts of distinct nodes are disjoint. Since T is rayless, the extended parts near-partition $\hat{V}(G)$. As an immediate consequence of $(*)$, every extended part of $(T, \mathcal{X})$ lives either in V_1 or V_2.

We colour the nodes of T using red and blue, as follows. Colour a node $t \in T$ red if $\hat{X}_t$ is non-empty and $\hat{X}_t \sqsubseteq V_1$. Similarly, we colour a node $t \in T$ blue if $\hat{X}_t$ is non-empty and $\hat{X}_t \sqsubseteq V_2$. Finally, we consider all the nodes $t \in \hat{T}$ with $\hat{X}_t = \emptyset$. These induce a forest in T. Colour all the nodes in a component of this forest red if the component has a red neighbour, and blue otherwise.

Let $T_1 \subseteq T$ be the forest induced by the red nodes, and let $T_2 \subseteq T$ be the forest induced by the blue nodes. The way in which we coloured the nodes with empty extended parts ensures that, for every connected component C of T_1 or of T_2, some node $t \in C$ has a non-empty extended part $\hat{X}_t$. Note that $B = E(\bigcup_{t \in T_1} X_t\,,\, \bigcup_{t \in T_2} X_t\,)$ by the definition of T_1 and T_2. We claim that we are done if T contains only finitely many T_1–T_2-edges. Indeed, if $s_0 t_0, \dots, s_n t_n$ are the finitely many T_1–T_2-edges with $s_\ell \in T_1$ and $t_\ell \in T_2$, then

$$(V_1, V_2) = \bigwedge_{C \text{ a component of } T_2} \; \bigvee_{\ell \,:\, t_\ell \in C} \; \alpha(s_\ell, t_\ell)\;.$$

Thus, it remains to show that T contains only finitely many T_1–T_2-edges. For this, we consider the tree $\tilde{T}$ that arises from T by contracting every component of T_1 and every component of T_2 to a single node. Since T is rayless, so is $\tilde{T}$. By Kőnig's infinity lemma, it remains to show that $\tilde{T}$ is locally finite.

Suppose for a contradiction that $d \in \tilde{T}$ is a vertex that has some infinitely many neighbours c_n $(n \in \mathbb{N})$. Recall that the sets $Y_c := \bigcup_{t \in c} \hat{X}_t$, where c is a node of $\tilde{T}$, are non-empty. We choose one point $u_n \in Y_{c_n}$ for every $n \in \mathbb{N}$, and we apply the star-comb lemma in the connected side $G[V_i]$ of B where all sets Y_{c_n} live to the infinite set $U := \{\, u_n : n \in \mathbb{N}\,\}$. Then we cannot get a star, because the finite fundamental cuts of $(T, \mathcal{X})$ induced by its T_i–d-edges would force the centre vertex to lie in Y_d, contradicting the fact that Y_d lives in V_{3-i}. Therefore, the star-comb lemma must return a comb contained in $G[V_i]$ and attached to U. Without loss of generality, each u_n is a tooth of this comb.

Let us consider the end of G that contains the spine of the comb. This end is contained in a $(k+1)$-edge-block $\beta \sqsubseteq V_i$. And β in turn is included in a set Y_c where c is a component of T_i. Hence $c \neq d$. But then the fundamental cut of $(T, \mathcal{X})$ which corresponds to the T_i–d-edge on the c–d-path in T separates a tail of the comb from infinitely many u_n, which is a contradiction. $\square$

6. Trees of tangles from tangle–tree duality

In this chapter, which is based on [40] and joint work with Christian Elbracht and Jakob Kneip, we present a way of deducing tree-of-tangles theorems from the abstract tangle–tree duality theorem.

Already in the original work by Robertson and Seymour the theory of tangles has two major theorems: the tree-of-tangles theorem and the tangle–tree duality theorem. These two form the main pillars of tangle theory.

As we have seen, the first of these theorems allows one to distinguish all the tangles in a tree-like way, displaying their relative position in the underlying combinatorial structure. Recall one of the most abstract variants of the tree-of-tangles theorem:

Theorem 5.4 ([25, Theorem 6])**.** *Let $\vec{S}$ be a submodular separation system in some universe $\vec{U}$ of separations and let $\mathcal{P}$ a set of profiles of S. Then S contains a nested set that distinguishes $\mathcal{P}$.* ⇝ p. 60

The tangle–tree duality theorem, on the other hand, provides a tree-like dual object to tangles which, if no tangle exists, serves as a witness that there can be no tangle. In this chapter we demonstrate the versatility of the most abstract version of this duality theorem: we deduce Theorem 5.4 and some of its variations from the tangle–tree duality theorem, reducing the two pillars of abstract tangle theory to a single pillar.

In order to use tangle–tree duality to deduce tree-of-tangles theorems like Theorem 5.4, we exploit the generality of the most abstract version of the tangle–tree duality theorem, which reads as follows:

Theorem 6.1 (Tangle–tree duality theorem [28, Theorem 4.3])**.** *Let U be a universe containing a finite separation system $S \subseteq U$ and let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars such that $\mathcal{F}$ is standard for $\vec{S}$ and $\vec{S}$ is $\mathcal{F}$-separable. Then exactly one of the following statements holds:* ⇝ p. 134

- *there is an $\mathcal{F}$-tangle of S;*
- *there is an S-tree over $\mathcal{F}$.*

The strength of Theorem 6.1 lies in the flexibility it allows in the choice of $\mathcal{F}$. This set $\mathcal{F}$ can be tailored to capture a wide variety of tangles and clusters, allowing Theorem 6.1 to be employed in a multitude of different settings ([25, 27]). The freedom in choosing and manipulating $\mathcal{F}$ will also allow us to achieve our goal of deducing tree-of-tangles theorems from Theorem 6.1: by a clever choice of $\mathcal{F}$ we can ensure that there is no $\mathcal{F}$-tangle of S, and that the S-tree over $\mathcal{F}$ one then obtains will be a tree of tangles. We will present multiple variations of this idea throughout the chapter.

In terms of simplicity and brevity, reducing the tree-of-tangles theorem to the tangle–tree duality theorem in this way cannot compete with its direct proofs, not with those in [25, 26] and certainly not the one in Chapter 5. Instead of competing in terms of simplicity and brevity just for a proof of the tree-of-tangles theorem, our aim here is to bridge the two parts of the theory needed for their classical proofs. This can be viewed in two ways.

Firstly, that we introduce tools from tangle–tree duality into the world of trees of tangles, which gives us a new method for building trees in this context very unlike the proofs in [25, 26] and Chapter 5. Secondly, and perhaps more importantly, from the perspective of tangle–tree duality this may be viewed as introducing a new range of ways of how to apply the duality theorem by a careful choice of $\mathcal{F}$. Previous applications of Theorem 6.1 all worked with largely similar choices of $\mathcal{F}$, all designed to capture some notion of 'width', whereas we specifically construct $\mathcal{F}$ in such a way that no $\mathcal{F}$-tangle can exist, thereby making sure that Theorem 6.1 gives us the dual object which will be the desired tree-of-tangles.

A new result that we get from this method is that it allows us to bound the degrees of the nodes in a tree of tangles in some contexts. Getting such a degree condition out of the original proofs does not appear to be simple.

To achieve our last result, we use a strengthened version of Theorem 6.1, the proof of which can be found in the paper [40] upon which this chapter is based or in Kneip's thesis [58].

In Section 6.1 we will introduce the required terminology of tangle–tree duality. In Section 6.2 we prove our first basic tree-of-tangles theorem, for structurally submodular separation systems. A refined version of this argument will be given in Section 6.3, where we show that the approach via tangle–tree duality yields a bound on the degrees of the nodes in a tree of tangles. In Section 6.4 we present a more involved argument to obtain a tree of tangles that distinguishes a set of profiles *efficiently*. Again, this approach can be used to obtain a result about the degrees in such a tree, and we do so in Section 6.5. In Section 6.6 we prove a tree-of-tangles theorem for tangles of different orders. In our final section, Section 6.6.1, we introduce the strengthened tangle–tree duality theorem and use it to obtain a tree-of-tangles theorem for profiles of different order.

6.1. Tangle–tree duality: terminology and tools

Since we combine the theory of tangle–tree duality and of trees of tangles we need the terminology of both. Consequently, this results in a large number of definitions which need to be understood for the comprehension of our proofs. In addition to the frameworks of [20, 25, 26] which are comprehensively used throughout the thesis, we further need the tools of tangle–tree duality from [28]. We will summarise the needed facts and terminology briefly, for an in-depth motivation of the set-up see [28] itself.

We will again be working with $\mathcal{F}$-tangles, particularly $\mathcal{F}$-tangles in submodular universes, and profiles. Observe that profiles are $\mathcal{P}$-tangles for the set $\mathcal{P}$ of all 'profile triples' $\{\vec{r}, \vec{s}, (\vec{r} \vee \vec{s})^*\} \subseteq \vec{S}$.

For the duality theorem however, we will consider sets $\mathcal{F}$ of *stars*: sets $\sigma \subseteq \vec{S}$ of non-degenerate separations such that $\vec{s} \leqslant \overleftarrow{t}$ for all $\vec{s}, \vec{t} \in \sigma$.

We say that a set $\mathcal{F}$ *forces* a separation $\vec{s} \in \vec{S}$ if $\{\overleftarrow{s}\} \in \mathcal{F}$. A set $\mathcal{F}$ is *standard for* $\vec{S}$ if it forces all trivial separations in $\vec{S}$, that is $\mathcal{F}$ contains all singletons $\{\overleftarrow{s}\}$ for $\vec{s} \in \vec{S}$ which are cotrivial in $\vec{S}$.

Recall that, given a separation system $\vec{S}$, an S-tree (T, α) is a tree T together with a function $\alpha\colon \vec{E}(T) \to \vec{S}$ which commutes with *, i.e., $\alpha(\overleftarrow{e}) = \alpha(\vec{e})^*$. The S-tree is order-respecting if α preserves the partial order: $\alpha(\vec{e}) \leqslant \alpha(\vec{f})$ whenever $\vec{e} \leqslant \vec{f}$, i.e., α is a homomorphism of separation systems.

Now, for $t \in V(t)$ we denote as $\alpha(t)$ the set $\{\alpha(st) : s \in N(t)\}$. Given some set $\mathcal{F}$ of subsets of S, an S-tree (T, α) is *over* $\mathcal{F}$ if $\alpha(t) \in \mathcal{F}$ for all $t \in V(T)$.

An S-tree (T, α) is *irredundant*, if for every node $t \in V(T)$ and distinct neighbours $t', t'' \in N(t)$ we have that $\alpha(t', t) \neq \alpha(t'', t)$.

Note that, if $\mathcal{F}$ is a set of stars, then every irredundant S-tree over $\mathcal{F}$ is order-respecting.

Given a separation system $\vec{S}$ inside a universe $\vec{U}$ and $\vec{r}, \vec{s}_0 \in \vec{S}$ with $\vec{s}_0 \geqslant \vec{r}$ and where $\vec{r}$ is neither degenerate nor trivial in $\vec{S}$, the *shifting map* $f{\downarrow}^{\vec{r}}_{\vec{s}_0}$ is defined by letting

$$f{\downarrow}^{\vec{r}}_{\vec{s}_0}(\vec{s}) = \vec{s} \vee \vec{s}_0 \quad \text{and} \quad f{\downarrow}^{\vec{r}}_{\vec{s}_0}(\overleftarrow{s}) = (\vec{s} \vee \vec{s}_0)^*$$

for every $\vec{s} \in \vec{S}_{\geqslant \vec{r}} \setminus \{\overleftarrow{r}\}$, where $S_{\geqslant \vec{r}}$ is the set of all separations $s \in S$ which have an orientation $\vec{s}$ with $\vec{s} \geqslant \vec{r}$, and $\vec{S}_{\geqslant \vec{r}}$ is the set of all orientations of separations in $S_{\geqslant \vec{r}}$.

For a leaf x of an irredundant S-tree (T, α) over some set of stars with $\{\overleftarrow{r}\} = \alpha(x)$, we write

$$\alpha_{x, \vec{s}_0} := f{\downarrow}^{\vec{r}}_{\vec{s}_0} \circ \alpha\,.$$

The resulting new tree $(T, \alpha_{x, \vec{s}_0})$ is called *the shift of* (T, α) *from* $\vec{r}$ *to* $\vec{s}_0$ if the leaf x is the only one which has $\alpha(x) = \{\overleftarrow{r}\}$.

Given a separation system $\vec{S}$ inside a universe $\vec{U}$ and a star $\sigma \subseteq \vec{S}$ a *shift of* σ *(to some* $\vec{s}_0 \in \vec{S}$*)* is a star of the form

$$\sigma^{\vec{s}_0}_{\vec{x}} := \{\vec{x} \vee \vec{s}_0\} \cup \{\vec{y} \wedge \overleftarrow{s}_0 : \vec{y} \in \sigma \setminus \{\vec{x}\}\},$$

where $\vec{x} \in \sigma$. Note that if, for some $\vec{r} \in \vec{S}$, we have $\vec{x} \geqslant \vec{r}$, then $\sigma^{\vec{s}_0}_{\vec{x}}$ is the image of σ under $f{\downarrow}^{\vec{r}}_{\vec{s}_0}$.

A separation $\vec{s}$ *emulates* $\vec{r}$ *in* $\vec{S}$ if $\vec{s} \geqslant \vec{r}$ and for every $\vec{t} \in \vec{S} \setminus \{\overleftarrow{r}\}$ with $\vec{t} \geqslant \vec{r}$ we have $\vec{s} \vee \vec{r} \in \vec{S}$. The separation $\vec{s}$ emulates $\vec{t}$ in $\vec{S}$ *for* $\mathcal{F}$ if additionally for every star $\sigma \in \mathcal{F}$ with $\overleftarrow{r} \notin \sigma$ and every $\vec{x} \in \sigma$ with $\vec{x} \geqslant \vec{r}$ we have $\sigma^{\vec{s}}_{\vec{x}} \in \mathcal{F}$.

Note that for an irredundant S-tree (T, α) over some set of stars $\mathcal{F}$ with $\{\overleftarrow{r}\} = \alpha(x)$, for some leaf x of T, the shift from $\vec{r}$ to $\vec{s}_0$ is again an S-tree over $\mathcal{F}$ if $\vec{s}_0$ emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$.

A separations system S is *separable* if for any two non-trivial non-degenerate separations $\vec{r}_1, \vec{r}_2 \in \vec{S}$ with $\vec{r}_1 \leqslant \overleftarrow{r}_2$ there exists a separation $\vec{s}_0 \in \vec{S}$, with $\vec{r}_1 \leqslant \vec{s}_0 \leqslant \overleftarrow{r}_2$ such that $\vec{s}_0$ emulates $\vec{r}_1$ in S and $\overleftarrow{s}_0$ emulates $\vec{r}_2$ in S. The separation system S is *$\mathcal{F}$-separable* if we can choose, for any two such $\vec{r}_1$ and $\vec{r}_2$ which are non-trivial non-degenerate and not forced by $\mathcal{F}$, such an $\vec{s}_0$ so that $\vec{s}_0$ emulates $\vec{r}_1$ in S for $\mathcal{F}$ and $\overleftarrow{s}_0$ emulates $\vec{r}_2$ in S for $\mathcal{F}$.

The abstract tangle–tree duality theorem now states the following:

Theorem 6.1 (Tangle–tree duality theorem [28, Theorem 4.3])**.** *Let U be a universe containing a finite separation system $S \subseteq U$ and let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars such that $\mathcal{F}$ is standard for $\vec{S}$ and $\vec{S}$ is $\mathcal{F}$-separable. Then exactly one of the following statements holds:*

- *there is an $\mathcal{F}$-tangle of S;*
- *there is an S-tree over $\mathcal{F}$.*

If, in the following, we speak of *the* duality theorem, we mean Theorem 6.1.

The condition of $\mathcal{F}$-separability is sometimes split into two parts which, in sum, are stronger: Firstly, that S is separable and secondly that $\mathcal{F}$ is *closed under shifting*, that is, every shift σ' of a star $\sigma \in \mathcal{F}$ is also in $\mathcal{F}$ if $\sigma' \subseteq \vec{S}$. (Compare [25, Lemma 12].)

We shall need the following additional lemmas from the literature:

Lemma 6.2 ([28, Lemma 2.1])**.** *Every irredundant S-tree (T, α) over stars is order-respecting. In particular, $\alpha(\vec{E}(T))$ is a nested set of separations in $\vec{S}$.*

Lemma 6.3 ([28, Lemma 2.2])**.** *Let (T, α) be an irredundant S-tree over a set $\mathcal{F}$ of stars. Let e, f be distinct edges of T with orientations $\vec{e} < \vec{f}$ such that $\alpha(\vec{e}) = \alpha(\overleftarrow{f}) =: \vec{r}$. Then $\vec{r}$ is trivial.*

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

Lemma 6.4 ([28, Lemma 2.3])**.** *If (T, α) is an S-tree over $\mathcal{F}$, possibly redundant, then T has a subtree T' such that (T', α') is an irredundant S-tree over $\mathcal{F}$, where α' is the restriction of α to $\vec{E}(T')$. If (T, α) is rooted at a leaf x and T has an edge, then T' can be chosen so as to contain x and e_x, the edge incident to x in T.*

Lemma 6.5 ([28, Lemma 2.4])**.** *Let (T, α) be an S-tree over a set $\mathcal{F}$ of stars, rooted at a leaf x. Assume that T has an edge, and that $\vec{r} = \alpha(\vec{e}_x)$ is non-trivial. Then T*

has a minor T' containing x and e_x such that (T', α'), where $\alpha' = \alpha \restriction \vec{E}(T')$, is a tight and irredundant S-tree over $\mathcal{F}$.

For every such (T', α') the edge $\vec{e}_x$ is the only edge $\vec{e} \in \vec{E}(T')$ with $\alpha(\vec{e}) = \vec{r}$.

Lemma 6.6 ([25, Lemma 13]). *Let $\vec{U}$ be a universe of separations and $\vec{S} \subseteq \vec{U}$ a structurally submodular separation system. Then $\vec{S}$ is separable.*

Moreover, we shall need a variant of [28, Lemma 4.2] which follows with the exact same proof:

Lemma 6.7 ([28]). *Let $\mathcal{F} \subseteq 2^{\vec{U}}$ be a set of stars. Let (T, α) be a tight and irredundant S-tree with at least one edge, over some set of stars, and rooted at a leaf x. Assume that $\vec{r} := \alpha(\vec{e}_x)$ is non-trivial and non-degenerate, let $\vec{s}_0 \in \vec{S}$ emulate $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$, and consider $\alpha' := \alpha_{x,\vec{s}_0}$. Then (T, α') is an order-respecting S-tree in which $\{ \overleftarrow{s}_0 \}$ is a star associated with x but with no other leaf of T. Moreover $\alpha'(t) \in \mathcal{F}$ for all $t \neq x$ with $\alpha(t) \in \mathcal{F}$.*

The only difference in the statement between Lemma 6.7 and [28, Lemma 4.2] is that [28, Lemma 4.2] requires that (T, α) is an S-tree *over $\mathcal{F}$*, whereas we only require (T, α) to be an S-tree *over some set of stars*. Consequently, in [28, Lemma 4.2] it is shown that then (T, α') is an S-tree over $\mathcal{F} \cup \{ \{ \overleftarrow{s}_0 \} \}$ whereas we only conclude that $\alpha'(t) \in \mathcal{F}$ whenever $\alpha(t) \in \mathcal{F}$.

6.1.1. Splices in submodular universes

In addition to the existing terminology, we shall use the following new concept, which has already been considered in [27], but was not given a name there: In a submodular universe $\vec{U}$ a separation $\vec{s}$ is a *splice for* a separation $\vec{r}$ with $\vec{r} \leqslant \vec{s}$ if there is no separation $\vec{t}$ with $\vec{r} \leqslant \vec{t} \leqslant \vec{s}$ and $|t| < |s|$. A *splice between* two separations $\vec{r}$ and $\vec{s}$ with $\vec{r} \leqslant \vec{s}$ is one of minimum order among all $\vec{t}$ with $\vec{r} \leqslant \vec{t} \leqslant \vec{s}$.

These splices are good choices for proving separability due to the next lemma. It follows directly from the proof of Lemma 3.4 of [27] which, phrased in our terminology, considers a splice between two separations. We recapitulate the main argument of this proof below.

Lemma 6.8 ([27]). *Consider $\vec{S_k} \subseteq \vec{U}$ in a submodular universe. If $\vec{s} \in \vec{S_k}$ is a splice for $\vec{r} \in \vec{S_k}$, then, for every $\vec{t} \in \vec{U}$ with $\vec{t} \geqslant \vec{r}$, the order of $\vec{t} \vee \vec{s}$ is at most the order of $\vec{t}$. In particular, $\vec{s}$ emulates $\vec{r}$ in $\vec{S_k}$.*

Proof sketch, see [27, Lemma 3.4]. If the order of $\vec{t} \vee \vec{s}$ were greater than the order of $\vec{t}$, then, by submodularity, the order of $\vec{t} \wedge \vec{s}$ would be less than the order of $\vec{s}$. However, by the fish Lemma 2.1, $\vec{r} \leqslant \vec{t} \wedge \vec{s} \leqslant \vec{s}$ and this contradicts the fact that $\vec{s}$ is a splice for $\vec{r}$. $\square$

This lemma then directly implies the ultimate statement of [27, Lemma 3.4]:

Lemma 6.9 ([27, Lemma 3.4]). *Every $\vec{S_k} \subseteq \vec{U}$ in a submodular universe is separable.*

6.2. Trees of tangles in submodular separation systems

In this section we will prove a first tree-of-tangles theorem. It is a theorem for regular profiles, all of the same structurally submodular separation system, and states as follows:

Theorem 6.10. *Let $\vec{S}$ be a submodular separation system in some universe of separations $\vec{U}$. Then S contains a nested set that distinguishes the set of regular profiles of S.*

By itself Theorem 6.10 is nothing special; indeed, it is a slight weakening of Theorem 5.4, which asserts the same but without requiring the profiles to be regular. In this case the ingredients of the proof are more interesting than its result: we shall obtain Theorem 6.10 as a direct consequence of Theorem 6.1.

So let $\vec{S}$ be a structurally submodular separation system inside some universe $\vec{U}$. Since we are interested in the regular profiles of S we may assume that S has no degenerate elements. Our strategy will be as follows: we shall construct a set $\mathcal{F} \subseteq 2^{\vec{U}}$ for which there is no $\mathcal{F}$-tangle of S, and so that every element of $\mathcal{F}$ is included in at most one regular profile of S. If we can achieve this, then Theorem 6.1 applied to this set $\mathcal{F}$ will yield an S-tree over $\mathcal{F}$. The set N of edge labels of this S-tree (T, α) will then be the desired nested set distinguishing all regular profiles of S: each regular profile P of S orients the edges of T and hence includes a star σ of the form $\alpha(t)$ for some $t \in V(T)$. By choice of $\mathcal{F}$, this σ is included in no other regular profile of S, which means that it distinguishes P from all other profiles.

To construct this set $\mathcal{F}$, first let $\mathcal{P}$ be the set of all 'profile triples' in $\vec{S}$: the set of all $\{\vec{r}, \vec{s}, (\vec{r} \vee \vec{s})^*\} \subseteq \vec{S}$. For a consistent orientation of S it is then equivalent to be a profile of S and to be a $\mathcal{P}$-tangle. Furthermore, let $\mathcal{C}$ be the set of all $\{\vec{s}\}$ with $\vec{s} \in \vec{S}$ co-small. Finally, let $\mathcal{M}$ consist of each of the sets $\max P$ of maximal elements of P for each regular profile P of S. We then take

$$\mathcal{F} := \mathcal{P} \cup \mathcal{C} \cup \mathcal{M}\,.$$

With these definitions the regular profiles of S are precisely its $(\mathcal{P} \cup \mathcal{C})$-tangles; and there are no $\mathcal{F}$-tangles of S since each regular profile P of S includes $\max P \in \mathcal{M} \subseteq \mathcal{F}$. If this $\mathcal{F}$ were a set of of stars and if we could feed this $\mathcal{F}$ to Theorem 6.1, we would receive an S-tree over $\mathcal{F}$ and the edge labels of this S-tree would be our desired nested set, since each element of $\mathcal{F}$ in included in at most one regular profile of S: indeed, the regular profiles of S have no subsets in $\mathcal{P}$ or $\mathcal{C}$, and each element $\max P \in \mathcal{M}$ in included only in P itself.

Unfortunately, we are still some way off from plugging $\mathcal{F}$ into Theorem 6.1: we need to ensure that $\mathcal{F}$ is a set of stars that is standard for S and that S is $\mathcal{F}$-separable. Out of these the second and one half of the third are easy: $\mathcal{F}$ is standard for S since $\mathcal{C} \subseteq \mathcal{F}$ is, and S is separable by Lemma 6.6.

We thus need to show that S is not only separable but $\mathcal{F}$-separable. While our current set $\mathcal{F}$ is not even a set of stars yet, in [23] a solution was laid out for this exact situation: a series of lemmas from [23] shows that we can simply *transform $\mathcal{F}$ into* a set of stars and close it under shifting without altering the set of $\mathcal{F}$-tangles of S.

The way to do this is as follows. Given two elements $\vec{r}$ and $\vec{s}$ of some set $\sigma \subseteq \vec{S}$, by submodularity, either $\vec{r} \wedge \overleftarrow{s}$ or $\overleftarrow{r} \wedge \vec{s}$ must lie in $\vec{S}$. *Uncrossing $\vec{r}$ and $\vec{s}$ in σ* then means to replace either $\vec{r}$ by $\vec{r} \wedge \overleftarrow{s}$ or $\vec{s}$ by $\overleftarrow{r} \wedge \vec{s}$, depending on which of these two lies in $\vec{S}$. (Structural submodularity ensures that at least one of them does.) Uncrossing all pairs of elements of σ in turn yields a star σ^*, which we call an *uncrossing* of σ. (Note that σ^* is not in general unique since it depends on the order in which one uncrosses the elements of σ.) It is then easy to see that a regular profile of S includes σ if and only if it includes σ^*:

Lemma 6.11 ([23, Lemma 11]). *If a regular profile of S includes an uncrossing of some set, it also includes that set.*

Conversely, if a regular consistent orientation of S includes some set, it also includes each uncrossing of that set.

Let us write $\mathcal{F}^*$ for the set of all uncrossings of elements of $\mathcal{F}$. Then $\mathcal{F}^*$ is a set of stars that is standard for S. We are still not done, however, since $\mathcal{F}^*$ need not be closed under shifting. We can fix this in a similar manner though.

Just as for uncrossings it is not hard to show that the inclusion of a star's shift in a regular profile implies that star's inclusion:

Lemma 6.12 ([23, Lemma 13]). *If a regular profile of S includes a shift of some star, it also includes that star.*

In [23] the definition of a shift of a star contains additional technical assumptions on σ and $\vec{s}_0$, keeping in line with the precise assumptions of Theorem 6.1. However the proof of Lemma 6.12 does not necessitate this, and neither does its application.

Lemma 6.12 says that if we close $\mathcal{F}^*$ under shifting we, again, do not alter the set of $\mathcal{F}^*$-tangles of S. Formally, set $\mathcal{G}_0 = \mathcal{F}^*$, and for $i \geqslant 1$ let $\mathcal{G}_i$ be the set of all shifts of stars in $\mathcal{G}_{i-1}$. We write $\widehat{\mathcal{F}^*} := \bigcup_{i \in \mathbb{N}} \mathcal{G}_i$. Then, by Lemma 6.12, the $\widehat{\mathcal{F}^*}$-tangles of S are precisely its $\mathcal{F}^*$-tangles, which is to say that there are no $\widehat{\mathcal{F}^*}$-tangles of S. Moreover, this set $\widehat{\mathcal{F}^*}$ still has the property that each star in it is included in at most one regular profile: let us say that $\widehat{\sigma^*} \in \widehat{\mathcal{F}^*}$ *originates from* $\sigma \in \mathcal{F}$ if $\widehat{\sigma^*}$ can be obtained by a series of shifts from an uncrossing of σ. Lemma 6.11 and 6.12 then say that if $\widehat{\sigma^*} \subseteq P$ for a regular profile P, and $\widehat{\sigma^*}$ originates from $\sigma \in \mathcal{F}$, then $\sigma \subseteq P$. Since the only element of $\mathcal{F}$ which P includes is $\max P$, this implies that no other regular profile of S includes $\widehat{\sigma^*}$.

We can thus formally prove Theorem 6.10:

Proof of Theorem 6.10. Define $\mathcal{P}$, $\mathcal{C}$, $\mathcal{M}$, $\mathcal{F}$, $\mathcal{F}^*$, and $\hat{\mathcal{F}}^*$ as above. Then $\hat{\mathcal{F}}^*$ is standard for S since $\mathcal{C} \subseteq \hat{\mathcal{F}}^*$, and closed under shifting by construction. By Lemma 6.6 S is separable. Together this gives that S is $\mathcal{F}$-separable. Hence we can apply the tangle–tree duality theorem, Theorem 6.1, to obtain either an $\hat{\mathcal{F}}^*$-tangle of S or an S-tree over $\hat{\mathcal{F}}^*$.

We claim that the first is impossible. For suppose that P is some $\hat{\mathcal{F}}^*$-tangle of S. From $\mathcal{C} \subseteq \hat{\mathcal{F}}^*$ we know that P is a regular and consistent orientation of S. If P has the profile property (P), then we could derive a contradiction from Lemma 6.11 and 6.12 since S has no $\mathcal{F}$-tangle. On the other hand, if P is not a profile, then P includes some set $\sigma \in \mathcal{P}$. By the second part of Lemma 6.11 P then also includes some (in fact: each) uncrossing of σ and hence a set in $\mathcal{F}^* \subseteq \hat{\mathcal{F}}^*$, contrary to its status as an $\hat{\mathcal{F}}^*$-tangle.

So let (T, α) be the S-tree over $\hat{\mathcal{F}}^*$ returned by Theorem 6.1, which we may assume to be irredundant (Lemma 6.4). Let $\vec{N}$ be the image of α. Then N is a nested subset of S (Lemma 6.2). Let us show that N distinguishes all regular profiles of S. Since (T, α) is an S-tree over $\hat{\mathcal{F}}^*$ each consistent orientation of S includes some star $\hat{\sigma^*} \in \hat{\mathcal{F}}^* \cap 2^{\vec{N}}$. In particular if P is a regular profile of S, then P includes some $\hat{\sigma^*} \in \hat{\mathcal{F}}^* \cap 2^{\vec{N}}$. Since the only element of $\mathcal{F}$ which P includes is $\max P$, this $\hat{\sigma^*}$ must originate from $\max P$. Consequently no other regular profile of S can include $\hat{\sigma^*}$, since none of them include $\max P$. Thus $\hat{\sigma^*}$ distinguishes P from every other regular profile of S. Since P was arbitrary this shows that N distinguishes all regular profiles of S. $\square$

Let us make some remarks on this proof of Theorem 6.10. First, in the definition of $\mathcal{F}$, we could have used other sets $\mathcal{M}$: the only properties of $\mathcal{M}$ that we used is that every regular profile of S contains some set from $\mathcal{M}$, and that no element of $\mathcal{M}$ is included in more than one such regular profile. We will put this observation to good use in Section 6.3, where we will make a more refined choice for $\mathcal{M}$ than simply collecting the sets of maximal elements from each profile.

Second, with the approach shown here it is not easy to strengthen Theorem 6.10 to the level of Theorem 5.4 by dropping the assumption of regularity, since Lemma 6.12 cannot do without this regularity.

In the remainder of this section we will show a more direct version of the proof presented above. This proof will be the guiding principle by which we will approach the issues of efficiency and profiles of differing order in Sections 6.4 and 6.6.

The core idea is that one can take as $\mathcal{F}$ the set of all stars that are included in at most one regular profile of S. An S-tree over this set $\mathcal{F}$ would immediately lead to the desired nested set distinguishing all regular profiles. Moreover this $\mathcal{F}$ is standard for S since $\mathcal{C} \subseteq \mathcal{F}$. To obtain this S-tree over $\mathcal{F}$ from Theorem 6.1 one would only need to show two things, namely that $\vec{S}$ is $\mathcal{F}$-separable and that there is no $\mathcal{F}$-tangle of S. The first of these amounts to Lemma 6.12; the second requires the two insights that every $\mathcal{F}$-avoiding consistent orientation is a regular profile, and that each regular profile of S includes some star in $\mathcal{F}$, both of which retrace some steps of Lemma 6.11.

Lemma 6.13. *Let $\vec{S} \subseteq \vec{U}$ be a submodular separation system in a universe $\vec{U}$ and let P be a profile of $\vec{S}$. There exists a star $\sigma \subseteq P$ such that no other profile of $\vec{S}$ includes σ.*

Proof. Let $\sigma \subseteq P$ be a star which minimizes the number of profiles which include σ. Suppose for a contradiction that there exists a profile $P' \neq P$ with $\sigma \subseteq P$. Some separation s, say, distinguishes P from P'. Clearly s crosses some element of σ.

Suppose that, subject to the above, σ and s are chosen so that the number of separations in σ that s crosses is minimum. Let $\vec{t} \in \sigma$ be a separation that s crosses. If either of the corner separations $\vec{t} \vee \vec{s}$ or $\vec{t} \vee \overleftarrow{s}$ was in $\vec{S}$, then, by the profile property, it would distinguish P and P'. It would also, by the fish Lemma 2.1, cross one less separation in σ than s does, contradicting the choice of s.

So by submodularity the corner separations $\vec{t} \wedge \vec{s}$ and $\vec{t} \wedge \overleftarrow{s}$ are in $\vec{S}$. Note that, by the profile property, any profile including

$$\sigma' := \sigma \setminus \{\, \vec{t} \,\} \cup \{\, \vec{t} \wedge \vec{s},\, \vec{t} \wedge \overleftarrow{s} \,\}$$

also includes σ. Consequently σ' together with s are a better choice than σ and s, a contradiction. $\square$

Lemma 6.14. *Given any set $\mathcal{P}$ of profiles of $\vec{S}$, every consistent orientation O of $\vec{S}$ which is not a profile in $\mathcal{P}$ contains a star σ which is not contained in any profile in $\mathcal{P}$.*

Proof. Since O is not a profile in $\mathcal{P}$ there is, for every profile P in $\mathcal{P}$, a separation s such that $\vec{s} \in O$ but $\overleftarrow{s} \in P$. Pick a set $\vec{N} \subseteq O$ which contains one such separation for every profile in $\mathcal{P}$ and is, subject to this, $\leqslant$-minimal: That is, there is no other such set N' together with an injective function $\alpha\colon N' \to N$ satisfying $\vec{s}' \leqslant \alpha(\vec{s}')$ for all $\vec{s}' \in N'$.

If N is a nested set, then N contains the desired star, so suppose that $\vec{s}, \vec{t} \in N$ cross. By submodularity we may suppose, after possibly renaming $\vec{s}$ and $\vec{t}$, that $\vec{s} \wedge \overleftarrow{t} \in S$ and thus, by consistency, $\vec{s} \wedge \overleftarrow{t} \in O$. We claim that $(N \setminus \{\, \vec{s} \,\}) \cup \{\, \vec{s} \wedge \overleftarrow{t} \,\}$ is also a candidate for N, contradicting the $\leqslant$-minimality. So suppose that $(N \setminus \{\, \vec{s} \,\}) \cup \{\, \vec{s} \wedge \overleftarrow{t} \,\}$ does not contain a separation $\vec{r}$ such that $\overleftarrow{r} \in P$, say. Then clearly $\overleftarrow{s} \in P$ and $\vec{t} \in P$, thus, by the profile property $\overleftarrow{s} \vee \vec{t} \in P$ which is precisely such an $\overleftarrow{r}$, a contradiction. $\square$

We are now ready to give a proof of Theorem 6.10 without resorting to Lemma 6.11:

Direct proof of Theorem 6.10. Let $\mathcal{P}$ be the set of regular profiles of S. Let $\mathcal{F}_{\mathcal{P}} \subseteq 2^{\vec{S}}$ consist of all stars $\sigma \subseteq \vec{S}$ for which one of the following is true:

(i) no profile in $\mathcal{P}$ includes σ, or

(ii) precisely one profile in $\mathcal{P}$ includes σ.

This $\mathcal{F}_{\mathcal{P}}$ is, by Lemma 6.12, closed under shifting: any shift of a star contained in at most one profile is again contained in at most one profile. The set $\mathcal{F}_{\mathcal{P}}$ is also standard for $\vec{S}$, since cosmall separations are contained in no regular profile.

By Theorem 6.1 there either exists an S-tree over $\mathcal{F}_{\mathcal{P}}$, or an $\mathcal{F}_{\mathcal{P}}$-tangle of S. In the former case we obtain the desired nested set. For the latter case observe that every $\mathcal{F}_{\mathcal{P}}$-tangle P, say, is a regular profile: By Lemma 6.14 every consistent orientation which avoids $\mathcal{F}_{\mathcal{P}}$ is a profile and if P would not be regular, it would contain a cosmall separation $\vec{s}$ which is impossible, since $\{\vec{s}\} \in \mathcal{F}_{\mathcal{P}}$. So by Lemma 6.13 there exists a star $\sigma \subseteq P$ which every profile other than P avoids. In particular $\sigma \in \mathcal{F}_{\mathcal{P}}$, which contradicts the fact that P is an $\mathcal{F}_{\mathcal{P}}$-tangle. $\square$

6.3. Application: Degrees in trees of tangles

In this section we are going to see that our proof of Theorem 6.10 in Section 6.2 has one advantage over the usual, more direct proofs of Theorem 6.10 from [25, 38]: It allows us to easily control the maximum degree of the resulting tree. More precisely: Let S be a submodular separation system in a universe U and P a regular profile of S. In this section we answer the following question: over all trees of tangles that distinguish all regular profiles of S, how low can the degree of the node containing P in those trees of tangles be?

Let us first make this notion of degree in a tree of tangles formal. For the purposes of this application only, a *tree of tangles (for S)* is an irredundant S-tree (T, α) whose set of edge labels distinguishes all regular profiles of S. For a regular profile P of S and a tree of tangles (T, α), the *node of P in T* is the unique sink of the orientation of T's edges induced by P, and the *degree of P in* (T, α) is the degree of this node.

Our question is thus: what is the minimum degree of P in (T, α) over all trees of tangles (T, α)?

A lower bound for this degree can be established as follows. Let $\delta(P)$ denote the minimal size of a set of separations which distinguishes P from all other regular profiles of S. If t is the node of P in some tree of tangles (T, α), then $\alpha(t)$ is such a set of separations which distinguishes P from all other regular profiles of S; thus, the degree of P in every tree of tangles (T, α) is at least $\delta(P)$.

We show that this lower bound can be achieved: there is a tree of tangles (T, α) for S in which P has degree exactly $\delta(P)$. In fact (T, α) will be optimal in this sense not just for P, but for all regular profiles of S simultaneously. Additionally the degrees of those nodes of (T, α) that are not the node of some regular profile will not be unreasonably high: the maximum degree of T will be attained in some profiles' node.

Theorem 27. *Let S be a submodular separation system in a universe of separations U. Then there is a tree of tangles (T,α) for S in which each regular profile P of S has degree exactly $\delta(P)$. Furthermore, if $\Delta(T) > 3$, then $\Delta(T) = \delta(P)$ for some regular profile P of S.*

To prove Theorem 27 we will follow the first proof of Theorem 6.10, making a more refined choice of $\mathcal{M}$, and utilize the fact that uncrossing and shifting a set cannot increase its size.

We will later see an example of a submodular separation system in which $\delta(P) \leqslant 2$ for every profile P but $\Delta(T) = 3$ for every tree of tangles T; this will demonstrate that the last assertion of Theorem 27 is optimal in that regard.

Observe further that the set of maximal elements of a profile P is a set which distinguishes P from every other profile of S. (In fact, the maximal elements of P distinguish P from every other consistent orientation of S.) Therefore $\delta(P) \leqslant |\max P|$ and hence the degree of P in the tree of tangles from Theorem 27 is at most $|\max P|$.

Proof of Theorem 27. For each regular profile P of S pick a subset $D_P \subseteq P$ of size $\delta(P)$ which distinguishes P from every other regular profile of S. Let $\mathcal{D}$ be the set of these D_P. Define $\mathcal{P}$ and $\mathcal{C}$ as in the proof of Theorem 6.10, and set

$$\mathcal{F} := \mathcal{P} \cup \mathcal{C} \cup \mathcal{D}\,.$$

From here, define $\mathcal{F}^*$ and $\hat{\mathcal{F}}^*$ just as in Theorem 6.10 and follow the same proof. The result is an S-tree over $\hat{\mathcal{F}}^*$, which we may assume to be irredundant and hence a tree of tangles for S.

Now let P be a regular profile of S, let t be the node of P in T, and $\hat{\sigma}^* := \alpha(t)$. As in the proof of Theorem 6.10 the only element of $\mathcal{F}$ from which $\hat{\sigma}^*$ can originate is D_P. Since uncrossing and shifting D_P cannot increase its size we have $|\hat{\sigma}^*| \leqslant |D_P| = \delta(P)$. Conversely we have $|\hat{\sigma}^*| \geqslant \delta(P)$ since $\hat{\sigma}^*$ distinguishes P from all other regular profiles. Thus the degree of P in (T,α) is indeed $\delta(P)$.

Finally, if $\Delta(T) > 3$, the maximum degree of T is attained in some node t whose associated star $\alpha(t)$ originates from some $D_P \in \mathcal{D}$, since all elements of $\hat{\mathcal{F}}^*$ originating from elements of $\mathcal{P}$ or $\mathcal{C}$ have size at most three. As above we thus have $|\alpha(t)| \leqslant |D_P| = \delta(P)$, giving $\Delta(T) = \delta(P)$. $\square$

Let us see an example showing that we cannot guarantee to find T with maximum degree less than three even if all regular profiles of S have $\delta(P) \leqslant 2$:

Example 6.15. Let V consist of the six points in Fig. 6.1, and S be the separation system given by the six outlined bipartitions of V together with $\{\emptyset, V\}$. The regular profiles of S correspond precisely to the six elements of V: each $v \in V$ induces a profile of S by orienting each bipartition towards v, and conversely each profile of S is of this form. Each profile P has at most two maximal elements, giving $\delta(P) \leqslant 2$. However, every tree of tangles for S must contain the outer three bipartitions and hence have a maximum degree of at least three.

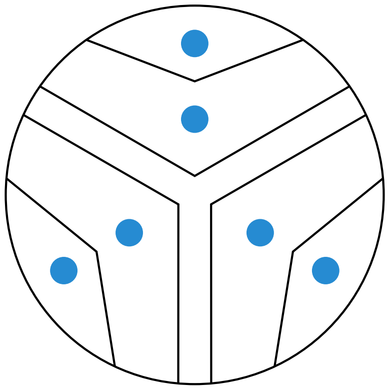

Figure 6.1.: A ground-set and system of bipartitions.

6.4. Efficient distinguishers

Often the submodularity of a separation system $\vec{S}$ is induced by a submodular order function from some submodular universe $\vec{U}$. In this case we are not just interested in a nested set of separations which distinguishes all profiles, but one which does so *efficiently*, that is, for any two profiles it contains a distinguishing separation of minimum possible order. In this section we are going to see how this can be achieved for regular profiles of a fixed $\vec{S}_k \subseteq \vec{U}$ utilising the duality theorem together with a separate application of its core mechanism: shifting S-trees.

We will prove this theorem:

Theorem 6.16. *Let $\vec{U}$ be a submodular universe and let $\mathcal{P}$ be a set of regular profiles of some $\vec{S}_k \subseteq \vec{U}$. Then there exists a nested set $N \subseteq S_k$ efficiently distinguishing all the profiles in $\mathcal{P}$.*

Our approach is similar to the one of the direct proof in Section 6.2, but we shall restrict our set of stars so that they do not interfere with efficiency.

Consider a nested set of separations which distinguishes all profiles efficiently and, subject to this, is $\subseteq$-minimal. Every profile P induces an orientation of this set, and the maximal elements of this orientation form a star. The separations in this star are, in a way, 'well connected' to the profile. We make this a condition on the stars we consider. For a star σ and a profile P, we say that σ *has the property* $\mathsf{Eff}(P)$ if the following holds:

$$\nexists\, \vec{s} \in \sigma \text{ and } \vec{s}\,' \in P\colon \vec{s} \leqslant \vec{s}\,' \text{ and } |s'| < |s|\,. \tag{$\mathsf{Eff}(P)$}$$

This condition ensures that, for two profiles P and P', a star σ with property $\mathsf{Eff}(P)$ containing $\vec{s}$, and a star σ' with property $\mathsf{Eff}(P')$ containing $\overleftarrow{s}$, the separation $\vec{s}$ needs to be an efficient P–P'-distinguisher. For if $\vec{s}$ is not efficient, consider an efficient P–P'-distinguisher $\vec{r} \in P$. Then $\vec{r}$ cannot be nested with $\vec{s}$, since $\vec{s} \leqslant \vec{r}$

would contradict property $\mathsf{Eff}(P)$ whereas $\vec{r} \leqslant \vec{s}$ would contradict property $\mathsf{Eff}(P')$. But $\vec{r}$ cannot cross $\vec{s}$ either: if it did we would have either $|\vec{r} \vee \vec{s}| < |\vec{s}|$ or $|\vec{r} \wedge \vec{s}| < |\vec{s}|$ by submodularity, again contradicting property $\mathsf{Eff}(P)$ or $\mathsf{Eff}(P')$, respectively.

Property $\mathsf{Eff}(P)$ is preserved under taking shifts:

Lemma 6.17. *Let $\vec{s} \in \vec{S}_k$ be a splice for $\vec{r} \in \vec{S}_k$ and let $\sigma \subseteq \vec{S}_k$ be a star with some $\vec{x} \in \sigma$ with $\vec{x} \geqslant \vec{r}$. If a profile P contains both σ and $\sigma' := \sigma^{\vec{s}}_{\vec{x}}$ and σ has property $\mathsf{Eff}(P)$, then also σ' has property $\mathsf{Eff}(P)$.*

Proof. Suppose for a contradiction that σ' does not have property $\mathsf{Eff}(P)$, that is, above some $\vec{t} \wedge \overleftarrow{s} \in \sigma'$, where $\vec{t} \in \sigma$, there is a separation $\vec{t}' \in P$ of lower order than $\vec{t} \wedge \overleftarrow{s}$.

We will first show that we may assume $\vec{t}' \leqslant \vec{t}$. Since $\vec{s}$ is a splice for $\vec{r}$ we have $|\vec{s} \wedge \overleftarrow{t}| \geqslant |\vec{s}|$, and thus by submodularity $|\overleftarrow{s} \wedge \vec{t}| \leqslant |\vec{t}|$. So if $\vec{t}' > \vec{t}$, then this contradicts the assertion that σ has property $\mathsf{Eff}(P)$. If however $\vec{t}'$ crosses $\vec{t}$, then, by the profile property of P and property $\mathsf{Eff}(P)$ of σ, the supremum $\vec{t}' \vee \vec{t}$ has at least the order of $\vec{t}$. By submodularity then $\vec{t}' \wedge \vec{t}$ has at most the order of $\vec{t}'$. This is also a separation in P which is above $\vec{t} \wedge \overleftarrow{s}$ and of lower order than $\vec{t} \wedge \overleftarrow{s}$, so we may consider it instead.

Now, since $\vec{s}$ is a splice for $\vec{r}$ we have that $|\overleftarrow{t}' \wedge \vec{s}| \geqslant |\vec{s}|$, so by submodularity $\vec{t}' \wedge \overleftarrow{s}$ has at most the order of $\vec{t}'$. But this $\vec{t}' \wedge \overleftarrow{s}$ is the same as $\vec{t} \wedge \vec{s}$ since $\vec{t} \geqslant \vec{t}' \geqslant \vec{t} \wedge \overleftarrow{s}$. So we have $|\vec{t} \wedge \vec{s}| \leqslant |\vec{t}'|$, which contradicts the assumption that $|\vec{t}'| < |\vec{t} \wedge \vec{s}|$. $\square$

We define $\mathcal{F}_e$ as the set of all stars $\sigma \subseteq \vec{S}_k$ which are contained in at most one profile in $\mathcal{P}$ and which, if they are contained in a profile $P \in \mathcal{P}$, fulfil property $\mathsf{Eff}(P)$.

From Lemma 6.12 and 6.17 immediately we obtain the following corollary:

Corollary 6.18. *S_k is $\mathcal{F}_e$-separable.*

However, an S-tree over $\mathcal{F}_e$ does not necessarily give rise to an efficient distinguisher set for $\mathcal{P}$ because we make no assumptions on those stars which are not contained in any profile. Our proof of Theorem 6.16 will need to make additional arguments on why an efficient such tree exists.

It would be much more elegant if we could introduce a condition, similar to $\mathsf{Eff}(\cdot)$, on the stars which are in no profile, so as to guarantee that any S_k-tree over these stars is as desired. However, all possible such properties that the authors could come up with failed to give $\mathcal{F}$-separability and there is reason to believe that such a solution is not possible: The critical part in the proof of Theorem 6.16 will make a global argument, specifically that of two shifts of one separation one is an efficient distinguisher. Separability on the other hand is defined in terms of each individual shift of a star.

For this section's analogue of Lemma 6.13, we define the *fatness* of a star σ as the tuple $(n_{k-1}, n_{k-2}, \ldots, n_1, n_0)$, where n_i is the number of separations of order i in σ. We will consider the lexicographic order on the fatness of stars.

Lemma 6.19. *Given a set $\mathcal{P}$ of regular profiles of $\vec{S}_k$, every profile $P \in \mathcal{P}$ includes a star in $\mathcal{F}_e$.*

Proof. By Lemma 6.13 P includes a star which is contained only in P. Take such a star σ which has lexicographically minimal fatness and suppose for a contradiction that σ does not have property $\mathsf{Eff}(P)$. So take $\vec{s} \in \sigma$ and $\vec{r} \in P$ with $\vec{s} \leqslant \vec{r}$ and $|r| < |s|$. Among the possible choices for $\vec{r}$, let $\vec{r}$ be one which crosses as few separations in σ as possible. If $\vec{r}$ were nested with σ, then the maximal elements of $\sigma \cup \{\, \vec{r} \,\}$ would form a star of lower fatness, thus we may suppose that $\vec{r}$ crosses some $\vec{x} \in \sigma$.

By the choice of $\vec{r}$, the corner separations $\vec{r} \vee \vec{x}$ and $\vec{r} \wedge \overleftarrow{x}$ must have strictly higher order than $|r|$ since both are $\geqslant \vec{s}$. Thus, by submodularity, the corner separations $\vec{r} \wedge \vec{x}$ and $\overleftarrow{r} \wedge \vec{x}$ have strictly lower order than $|x|$. Now the star $\sigma' := \sigma \setminus \{\, \vec{x} \,\} \cup \{\, \vec{r} \wedge \vec{x}, \overleftarrow{r} \wedge \vec{x} \,\}$ has a lower fatness. This star is still contained in P by consistency and in no other profile, since every profile which includes σ' also includes σ by the profile property applied with $\vec{x}$ and $\vec{r}$. This contradicts the choice of σ. □

We are now able to prove Theorem 6.16:

Proof of Theorem 6.16. We may apply Theorem 6.1 for $\mathcal{F}_e$ since $\vec{S}_k$ is $\mathcal{F}_e$-separable by Corollary 6.18 and $\mathcal{F}_e$ is standard since cotrivial separations are not contained in any regular profile. From this theorem we cannot get an $\mathcal{F}_e$-tangle: such a tangle cannot be a profile in $\mathcal{P}$ by Lemma 6.19, and Lemma 6.14 states that every consistent orientation which is not a profile in $\mathcal{P}$ includes a star which is not contained in any profile in $\mathcal{P}$, but each of these stars is contained in $\mathcal{F}_e$, so no such orientation is an $\mathcal{F}_e$-tangle. So instead, there exists an S_k-tree over $\mathcal{F}_e$.

Among all S_k-trees over $\mathcal{F}_e$ pick an irredundant one, (T, α) say, whose associated separations efficiently distinguishes as many pairs of profiles as possible. Let us suppose that some pair of profiles P_1, P_2 is not distinguished efficiently by this tree.

Consider the nodes v_{P_1}, v_{P_2} of this tree corresponding to P_1 and P_2. These nodes are distinct, since every star in $\mathcal{F}_e$ is contained in at most one profile. Moreover, we can assume without loss of generality, that in no node on the path between v_{P_1} and v_{P_2} there lives a profile Q: In that case either the pair P_1, Q or the pair Q, P_2 would not be efficiently distinguished by (T, α) either, so we could consider them instead.

Let $\vec{s}_{P_1}$ be the separation associated to the first edge on the path from v_{P_1} to v_{P_2} and let $\vec{s}_{P_2}$ be the separation associated to the first edge on the path from v_{P_2} to v_{P_1}. There exists a separation t which efficiently distinguishes P_1 and P_2 and is nested with $\vec{s}_{P_1}$ and $\vec{s}_{P_2}$: if $\vec{t} \in P_1$ is not nested with, say $\vec{s}_{P_1}$, we know by property $\mathsf{Eff}(P)$ that $\overleftarrow{s}_{P_1} \vee \vec{t}$ needs to have order at least $|s_{P_1}|$, thus $\overleftarrow{s}_{P_1} \wedge \vec{t}$ has order at most $|t|$, so it efficiently distinguishes P_1 and P_2 and is nested with $\vec{s}_P$. Thus by the fish Lemma 2.1, there indeed needs to exists such a t which efficiently distinguishes P_1 and P_2 and is nested with $\vec{s}_{P_1}$ and $\vec{s}_{P_2}$. Moreover, t has an orientation such that

$\vec{s}_{P_1} \leqslant \vec{t} \leqslant \overleftarrow{s}_{P_2}$, otherwise the existence of t again contradicts either property $\mathsf{Eff}(P)$ or $\mathsf{Eff}(Q)$. Note that $\vec{t}$ thus is a splice between $\vec{s}_{P_1}$ and $\overleftarrow{s}_{P_2}$ and therefore $\vec{t}$ emulates $\vec{s}_{P_1}$ for $\mathcal{F}_e$ and $\overleftarrow{t}$ emulates $\vec{s}_{P_2}$ for $\mathcal{F}_e$.

Let T_{P_1} be the subtree of T consisting of the component of $T - v_{P_1}$ which contains v_{P_2} together with v_{P_1} and similarly let T_{P_2} be the subtree consisting of the component of $T - v_{P_2}$ containing v_{P_1} together with v_{P_2}.

We consider the trees (T_{P_1}, α_{P_1}) and (T_{P_2}, α_{P_2}) obtained from the restrictions $(T_{P_1}, \alpha \restriction T_{P_1})$ and $(T_{P_2}, \alpha \restriction T_{P_2})$ by applying the shifts $f\downarrow_{\vec{t}}^{\vec{s}_{P_1}}$ and $f\downarrow_{\overleftarrow{t}}^{\vec{s}_{P_2}}$, respectively. Consider now the tree (T', α') obtained from these two trees by identifying the respective edges associated with $\vec{t}$. By applying Lemma 6.7 with the two shifted trees the combined tree is again over $\mathcal{F}_e$. We may again assume it to be irredundant. We are going to show that it efficiently distinguishes more pairs of profiles than (T, α).

Let Q_1, Q_2 be a pair profiles which were efficiently distinguished by a separation $\vec{r}$ associated to an edge of (T, α). If $\vec{r}$ is not associated to any edge of (T', α'), then, without loss of generality, either $\vec{s}_{P_1} \leqslant \vec{r} \leqslant \vec{s}_{P_2}$ or both $\vec{s}_{P_1} \leqslant \vec{r}$ and $\vec{s}_{P_2} \leqslant \vec{r}$.

In the first case $\vec{r}$ distinguishes P_1 and P_2 and therefore $|\vec{r}| > |\vec{t}|$. By the definition of the shift, our tree (T', α') contains both, $\vec{r} \vee \vec{t}$ and $\vec{r} \wedge \vec{t}$, and both of them have order at most the order of $\vec{r}$, by Lemma 6.8. However, one of $\vec{r} \vee \vec{t}, \vec{r} \wedge \vec{t}$ and $\vec{t}$ distinguishes Q_1 and Q_2 and does so efficiently.

In the second case, by the definition of the shift, our tree (T', α') contains both, $\vec{r} \vee \vec{t}$ and $\vec{r} \vee \overleftarrow{t}$, and both of them have order at most the order of $\vec{r}$, again by Lemma 6.8. Again, one of $\vec{r} \vee \vec{t}$ and $\vec{r} \vee \overleftarrow{t}$ distinguishes Q_1 and Q_2 and does so efficiently.

Thus, since (T', α') additionally efficiently distinguishes P_1 and P_2 with t, this contradicts the choice of (T, α). $\square$

6.5. Degrees in efficient trees of tangles

In this section we apply our method from Section 6.3 to Theorem 6.16 to obtain a tree of tangles of low degree, but this time one which efficiently distinguishes the profiles. That is, we are interested in the minimal degrees of a tree of tangles whose associated separations efficiently distinguish all regular profiles of S_k.

Extending the definitions of Section 6.3, let us say that a tree of tangles (T, α) for S_k is *efficient*, if the set of edge labels not only distinguishes all regular profiles of S_k, but does so efficiently.

Given a k-profile P, we denote by $\delta_e(P)$ the minimal size of a star $\sigma \subseteq P$ with property $\mathsf{Eff}(P)$ which distinguishes P from all other regular profiles of S_k, i.e., every other regular profile orients some $\vec{s} \in \sigma$ as $\overleftarrow{s}$. Note that, by Lemma 6.19, there exists such a star for every regular profile P, thus $\delta_e(P)$ is a well defined natural number.

We denote by $\delta_{e,\mathsf{max}}$ the maximum of $\delta_e(P)$ over all regular profiles P.

We can give a bound on $\delta_e(P)$ which is not in terms of stars or nested sets:

Lemma 6.20. *Let P be a regular k-profile in U and let $D_P \subseteq P$ be a subset of P which contains, for every regular k-profile $P' \neq P$ in U, a separation which efficiently distinguishes P from P'. Let us denote as m the number of maximal elements of D_P. Then $\delta_e(P) \leqslant m$.*

Proof. It is enough to consider a set $D_P \subseteq P$ such that $m = |\max D_P|$ is as small as possible. Moreover, we may assume without loss of generality that every element of D_P distinguishes P efficiently from some other profile in $\mathcal{P}$, since we could otherwise remove it from D_P. We may furthermore assume that, subject to all this, D_P is chosen so that $\max D_P$ is $\leqslant$-minimal. Furthermore we may suppose that, for separations $\vec{r} \leqslant \vec{s}$ in D_P, the order of $\vec{r}$ is lower than the order of $\vec{s}$, since otherwise we could just remove $\vec{r}$ from D_P.

If the maximal separations in D_P are pairwise nested, they satisfy property $\mathsf{Eff}(P)$ by the fact that they distinguish P efficiently from some other profile P'. Further, every profile P' is distinguished from P by some maximal separation in D_P: there is an efficient P–P'-distinguisher $\vec{s} \in D_P$ and thus a maximal separation $\vec{t} \geqslant \vec{s}$ in D_P also distinguishes P from P'. Hence, if the maximal elements of D_P are pairwise nested, they are a candidate for $\delta_e(P)$ and therefore witness that $\delta_e(P) \leqslant m$.

So suppose that this is not the case, so two maximal separations $\vec{s}, \vec{t} \in D_P$ cross and, without loss of generality, $|\vec{s}| \leqslant |\vec{t}|$. By the definition of D_P, there is a profile P_s which is efficiently distinguished from P by $\vec{s} \in D_P$. Similarly, there is such a profile P_t for $\vec{t}$.

Since D_P was chosen to have as few maximal elements as possible, the separation $\vec{s} \vee \vec{t}$ has greater order than t: otherwise we could, by consistency and the profile property, replace $\vec{t}$ in D_P by $\vec{s} \vee \vec{t}$. Thus, by submodularity, the order of $\vec{s} \wedge \vec{t}$ is less than the order of $\vec{s}$. In particular, by efficiency of s and t, neither P_s nor P_t contains $(\vec{s} \wedge \vec{t})^* = \overleftarrow{s} \vee \overleftarrow{t}$.

Thus $\vec{s} \wedge \overleftarrow{t}$ and $\overleftarrow{s} \wedge \vec{t}$ have order precisely $|\vec{s}|$ and $|\vec{t}|$, respectively: if one of them had lower order this would, by the profile property, contradict the fact that s or t, respectively, efficiently distinguishes P from P_s or P_t, respectively. This means that, in particular, $\vec{s} \wedge \overleftarrow{t}$ efficiently distinguishes P from P_s.

For every $\vec{r} \leqslant \vec{s}$ in D_P we have assumed $|\vec{r}| < |\vec{s}|$. Both $\vec{r} \wedge \vec{t}$ and $\vec{r} \wedge \overleftarrow{t}$ have at most the order of $\vec{r}$ due to submodularity, the efficiency of $\vec{t}$, the profile property and consistency, analogue to the above.

Let us consider the set D'_P obtained from D_P by removing all $\vec{r} \leqslant \vec{s}$, and adding $\vec{s} \wedge \overleftarrow{t}$ as well as, for every $\vec{r} \leqslant \vec{s}$, any $\vec{r} \wedge \vec{t}$ and $\vec{r} \wedge \overleftarrow{t}$ which efficiently distinguishes P from some other profile. By the above, this set D'_P distinguishes P from every other regular profile, and is a candidate for D_P. The maximal separations of D'_P and of D_P are the same except that $\vec{s}$ in D_P is replaced by $\vec{s} \wedge \overleftarrow{t}$ in D'_P. This contradicts the choice of D_P with $\leqslant$-minimal maximal elements. $\square$

To limit the degree of the node of P in our tree of tangles we want to remove from $\mathcal{F}_e$ all the stars which are contained in P but are larger than $\delta_e(P)$. In order to achieve a maximum degree of $\delta_{e,\mathsf{max}}$ we also need to limit the size of the stars in $\mathcal{F}_e$ which are contained in no profile to $\delta_{e,\mathsf{max}}$. As in Section 6.3, we cannot limit the maximum degrees below 3. Along the lines of the proof of Lemma 6.11, the next lemma shows that we can find, in every consistent orientation O of $\vec{S}_k$ which is not a profile, a star of size 3 contained in O and in no profile.

Lemma 6.21. *Every consistent orientation O of $\vec{S}_k$ which is not a profile contains a star σ of size 3 which is not contained in any profile.*

Proof. As O is not a profile, there are $\vec{s}, \vec{t} \in O$ such that $\overleftarrow{s} \wedge \overleftarrow{t} \in O$. By submodularity, either $\vec{s} \wedge \overleftarrow{t}$ or $\overleftarrow{s} \wedge \vec{t} \in \vec{S}$, let us suppose the former one. Then $\sigma = \{\, \vec{s} \wedge \overleftarrow{t}, \vec{t}, \overleftarrow{s} \wedge \overleftarrow{t} \,\}$ is a star in O and σ cannot be contained in any profile: any profile P needs to contain either $\vec{s}$ or $\overleftarrow{s}$, and the profile property implies that P then cannot contain both, $\vec{s} \wedge \overleftarrow{t}$ and $\overleftarrow{s} \wedge \overleftarrow{t}$. □

We can now show the following variant of Theorem 6.16, which shows that we can find a tree of tangles of bounded degree:

Theorem 28. *Let $\vec{U}$ be a submodular universe and let $\mathcal{P}$ be the set of regular profiles of $\vec{S}_k$. Then there exists tree of tangles (T, α) such that, for every profile $P \in \mathcal{P}$, the degree of P in (T, α) is $\delta_e(P)$ and the maximal degree of T is at most $\mathsf{max}\{\, \delta_e(\mathcal{P}), 3 \,\}$.*

Proof. Let $\mathcal{F}^s_e$ be the subset of $\mathcal{F}_e$ consisting of, for every profile P, all stars from $\mathcal{F}_e$ of size $\delta_e(P)$ contained in P, together with all stars of size at most $\mathsf{max}\{\, \delta_e(\mathcal{P}), 3 \,\}$ from $\mathcal{F}_e$ not contained in any profile. For any star σ and any shift $\sigma^{\vec{r}}_{\vec{s}}$ of σ we have $|\sigma| \geqslant |\sigma^{\vec{r}}_{\vec{s}}|$. Further, S_k is $\mathcal{F}_e$-separable by Corollary 6.18. Moreover, the shift of a star cannot contain any profile which does not contain the original star by Lemma 6.12, thus S_k is also $\mathcal{F}^s_e$-separable.

Thus, all we need to show is that applying Theorem 6.1 cannot result in an $\mathcal{F}^s_e$-tangle, the rest of the proof can then be carried out as the proof of Theorem 6.16: Instead of S-trees over $\mathcal{F}_e$ we now consider S-trees over $\mathcal{F}^s_e$, and observe that the shifting argument in the proof of Theorem 6.16 again shifts stars in $\mathcal{F}^s_e$ to stars in $\mathcal{F}^s_e$.

However, applying Theorem 6.1 indeed cannot result in an $\mathcal{F}^s_e$-tangle: Such a tangle cannot be a regular profile, since by our definition of $\delta_e(P)$, there is a star in $\mathcal{F}^s_e$ contained in P. But every consistent orientation which is not a regular profile either contains a star $\{\vec{s}\}$ for a cosmall separation $\vec{s}$ – each such star is also contained in $\mathcal{F}_e$ – or contains, by Lemma 6.21 a star of size 3 not contained in any profile. Either such star is also contained in $\mathcal{F}^s_e$ by definition. □

6.6. Tangles of mixed orders

In this section we would like to use the ideas from Section 6.4 to obtain a proof of the following theorem by Diestel, Hundertmark, and Lemanczyk using tangle–tree duality:

Theorem 6.22 ([26, Corollary 3.7 without canonicity]). *Let $(\vec{U}, \leqslant, ^*, \vee, \wedge, |\cdot|)$ 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:*

1. *every two profiles in $\mathcal{P}$ are efficiently distinguished by some separation in T;*
2. *every separation in T efficiently distinguishes a pair of profiles in $\mathcal{P}$.*

Note that the original statement [26, Corollary 3.7] included a third property which guaranteed that the resulting set T is invariant under automorphisms. Our method, using the tangle–tree duality theorem, will not allow us to guarantee this, which is why we exclude the property in this version of [26, Corollary 3.7]. For more discussion of this property, *canonicity*, see [26, 34] as well as our proof of the full Corollary 3.7 of [26] in Section 5.5.

The challenge of Theorem 6.22 compared to Theorem 6.16 is that the set of profiles $\mathcal{P}$ considered in Theorem 6.22 consists of profiles of different orders. In particular, there might be profiles P_1 and P_2 in $\mathcal{P}$ which are efficiently distinguished by separations of order k, say, and there might be another profile $Q \in \mathcal{P}$ which has only order $l < k$ and thus does not orient the separations which efficiently distinguish P_1 and P_2. Thus, we cannot simply require the stars in our set $\mathcal{F}$ to be contained in at most one profile: the resulting S-tree over $\mathcal{F}$ would not necessarily distinguish all profiles in $\mathcal{P}$, for example it might not distinguish the profiles P_1 and Q from above. Our solution to this problem will be to restrict the set of stars further by additionally requiring that all the separations in a star in $\mathcal{F}$ 'could be oriented' by every profile in $\mathcal{P}$, even if that profile has lower order than the separation considered.

With this further restricted set of stars however S will no longer be $\mathcal{F}$-separable, but it will only fail to do so under rather specific circumstances. Thus in order to obtain a result in the fashion of Theorem 6.22, we shall use a slightly stronger version of Theorem 6.1, which allows us to exclude this specific situation im the requirement of $\mathcal{F}$-separability. The proof of this stronger version of Theorem 6.1 is due to Kneip and can be found in the paper version [40] of this chapter as well as in Kneip's thesis [58]. The statement of the strengthened duality theorem is this:

Theorem 6.23. *Let U be a finite universe, $S \subseteq U$ a separation system, and $\mathcal{F} \subseteq 2^{\vec{S}}$ a set of stars such that $\mathcal{F}$ is standard for S and S is critically $\mathcal{F}$-separable. Then precisely one of the following holds:*

- *there is an S-tree over $\mathcal{F}$;*
- *there is an $\mathcal{F}$-tangle of S.*

This theorem is a strengthening in the sense that it weakens the technical assumption that S be $\mathcal{F}$-separable to only require $\mathcal{F}$-separability for those separations whose inverse lies in no star of $\mathcal{F}$, rather than for all separations in $\vec{S}$. Formally:

A separation $\vec{r}$ in $\vec{S}$ is *$\mathcal{F}$-critical* if $\vec{r} \in \sigma$ for some $\sigma \in \mathcal{F}$, but there is no $\sigma' \in \mathcal{F}$ with $\sigma' \cap r = \{\overleftarrow{r}\}$. Observe that if $\vec{r} \in \vec{S}$ is $\mathcal{F}$-critical, then $\vec{r}$ is non-degenerate and not forced by $\mathcal{F}$, and in particular $\vec{r}$ is non-trivial in S since $\mathcal{F}$ is standard for S. We say that S is *critically $\mathcal{F}$-separable* if for all $\mathcal{F}$-critical $\vec{r}, \overleftarrow{r}' \in \vec{S}$ with $\vec{r} \leqslant \vec{r}'$ there exists an $s_0 \in S$ with an orientation $\vec{s}_0$ that emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$ and such that $\overleftarrow{s}_0$ emulates $\overleftarrow{r}'$ in $\vec{S}$ for $\mathcal{F}$. Clearly, if S is $\mathcal{F}$-separable, then S is critically $\mathcal{F}$-separable.

6.6.1. Obtaining a tree of mixed-order tangles from tangle–tree duality

Theorem 6.23 now allows us to use our methods from Theorem 6.16 to prove a tree-of-tangles theorem for different order tangles. More specifically we will obtain a result similar to Theorem 6.22, however our construction only works in distributive universes – that is, $\vec{r} \vee (\vec{s} \wedge \vec{t}) = (\vec{r} \vee \vec{s}) \wedge (\vec{r} \vee \vec{s})$, always – since we need the following result from [31], which is also found in [26]:

Lemma 6.24 ([26, Theorem 3.11], [31, Theorem 1], strong profile property). *Let $\vec{U}$ be a distributive universe and $\vec{S} \subseteq \vec{U}$ structurally submodular, then for any profile P of S and any $\vec{r}$ and $\vec{s} \in P$ there does not exists any $\vec{t} \in P$ such that $\overleftarrow{r} \vee \overleftarrow{s} \leqslant \vec{t}$.*

Moreover, our method will not allow us to distinguish all robust profiles, instead we need a slight strengthening of robustness: We say that a k-profile P is *strongly robust*, if for any $\vec{s} \in P$ and $\vec{r} \in \vec{U}$ where $\vec{s} \vee \vec{r}$ and $\vec{s} \vee \overleftarrow{r}$ both have at most the order of $\vec{s}$ one of $\vec{s} \vee \vec{r}$ and $\vec{s} \vee \overleftarrow{r}$ is in P. Note that most instances of tangles, for example tangles in graphs, are strongly robust profiles.

For this section let $\vec{U}$ be a distributive submodular universe and let $\mathcal{P}$ be some set of pairwise distinguishable strongly robust profiles in $\vec{U}$ (possibly of different order).

To handle the issue, that not all separations in a tree-of-tangles for profiles of different orders are oriented by all the considered profiles, we introduce the following additional definition: A consistent orientation O of S_k *weakly orients a separation s as $\vec{s}$* if O contains a separation $\vec{r}$ such that $\vec{s} \leqslant \vec{r}$. If we want to omit s we just say *O weakly contains $\vec{s}$.*

We will now only consider stars of separations where every separation is at least weakly oriented by all the profiles in $\mathcal{P}$. Specifically, we work with the set $\mathcal{F}_d$ consisting of all stars σ with the following properties:

1. There exists at most one profile $P \in \mathcal{P}$ such that $\sigma \subseteq P$.

2. For every profile $P \in \mathcal{P}$ such that $\sigma \nsubseteq P$ there exists $\vec{s} \in \sigma$ such that P weakly orients s as $\overleftarrow{s}$.

3. If there exists a $P \in \mathcal{P}$ such that $\sigma \subseteq P$, then σ satisfies property $\mathsf{Eff}(P)$.

We want to show that U is critically $\mathcal{F}_d$-separable, and our first step to do so is to show that splices – which we want to use in separability – are weakly oriented by every profile in $\mathcal{P}$.

Lemma 6.25. *Let $\vec{U}$ be a distributive submodular universe and let $\mathcal{P}$ be a set of strongly robust profiles in $\vec{U}$. Suppose that $\vec{r}$ and $\vec{s}$ are $\mathcal{F}_d$-critical separations in $\vec{U}$ with $\vec{r} \leqslant \overleftarrow{s}$, then every splice between $\vec{r}$ and $\overleftarrow{s}$ is weakly oriented by every profile in $\mathcal{P}$.*

Proof. Since $\vec{r}$ and $\vec{s}$ are $\mathcal{F}_d$-critical, they are contained in some star in $\mathcal{F}_d$ and hence weakly oriented by every profile in $\mathcal{P}$.

Let t be a splice between $\vec{r}$ and $\overleftarrow{s}$. If t is not weakly oriented by every profile in $\mathcal{P}$, then $\mathcal{P}$ contains a profile P of order at most $|t|$ which weakly orients r as $\vec{r}$ and s as $\vec{s}$, since every witnessing separation that a profile weakly orients r as $\overleftarrow{r}$ or s as $\overleftarrow{s}$ also witnesses that it weakly orients t. Let M_r^P be the set of all separations $\vec{w}_r$ in P satisfying $\vec{r} \leqslant \vec{w}_r$ and having minimal possible order with that property. Let $\vec{w}_r \in M_r^P$ be chosen $\leqslant$-maximally. Let $\vec{w}_s$ be defined for s, accordingly.

Observe that if $\vec{w}_r \leqslant \overleftarrow{s}$, respectively, then, by the order-minimality of M_r^P, the order of $\vec{w}_r$ is at least $|t|$ so P orients t, which contradicts the assumption that P does not weakly orient t. Similarly, $\vec{w}_s \leqslant \overleftarrow{r}$ results in a contradiction.

Suppose now that $\vec{w}_r$ crosses $\vec{s}$.

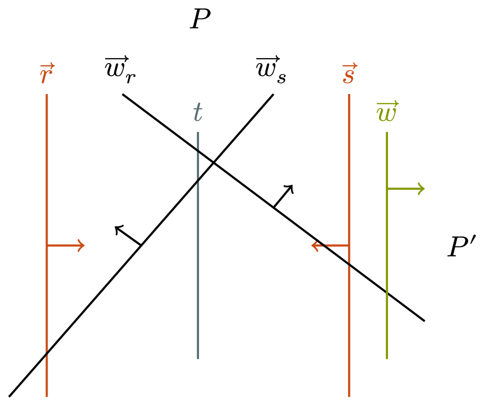

We claim that every profile P' in $\mathcal{P}$ which weakly orients s as $\overleftarrow{s}$ also weakly contains either $\overleftarrow{s} \vee \overleftarrow{w}_r$ or $\overleftarrow{s} \vee \vec{w}_r$. This then impies that $\{\, \overleftarrow{s}, \vec{s} \wedge \vec{w}_r, \vec{s} \wedge \overleftarrow{w}_r \,\}$ is a star in $\mathcal{F}_d$, which will contradict the $\mathcal{F}_d$-criticality of $\vec{s}$.

So suppose that P' weakly orients s as $\overleftarrow{s}$, witnessed by some $\vec{w} \in P'$ with $\vec{w} \geqslant \overleftarrow{s}$.

If $\vec{w}_r \vee \vec{w}$ had order at most the order of $\vec{w}_r$, this would contradict the choice of $\vec{w}_r$: By Lemma 6.24 applied to the separations $\vec{w}_r, \vec{w}_s, \overleftarrow{w} \wedge \overleftarrow{w}_r$, the profile P would need to contain $\vec{w}_r \vee \vec{w}$ which contradicts the choice of $\vec{w}_r$ being $\leqslant$-maximal in M_r^P.

Similarly, if $\overleftarrow{w} \wedge \vec{w}_r$ had order less than the order of $\vec{w}_r$, this would contradict the choice of $\vec{w}_r$: By consistency P would need to contain $\overleftarrow{w} \wedge \vec{w}_r$ which contradicts the definition of M_r^P, from which $\vec{w}_r$ was chosen.

Thus, by submodularity, $\vec{w} \wedge \vec{w}_r$ has order less than the order of w, and $\vec{w} \wedge \overleftarrow{w}_r$ has order at most the order of w. Hence, as P' is strongly robust, P' contains either $\overleftarrow{w} \vee \overleftarrow{w}_r$ or $\overleftarrow{w} \vee \vec{w}_r$ and therefore either weakly orients $\vec{s} \wedge \vec{w}_r$ as $\overleftarrow{s} \vee \overleftarrow{w}_r$ or $\vec{s} \wedge \overleftarrow{w}_r$ as $\overleftarrow{s} \vee \vec{w}_r$.

This proves the claim which results in a contradiction to the assumption that $\vec{s}$ is $\mathcal{F}_d$ critical. Thus we may suppose that $\vec{w}_r$ does not cross $\vec{s}$ and, by a symmetric argument, that $\vec{w}_s$ does not cross $\vec{r}$. Hence $\vec{r} \leqslant \vec{w}_s$ and $\vec{s} \leqslant \vec{w}_r$. We may therefore assume without loss of generality that $\vec{w}_r = \vec{w}_s$.

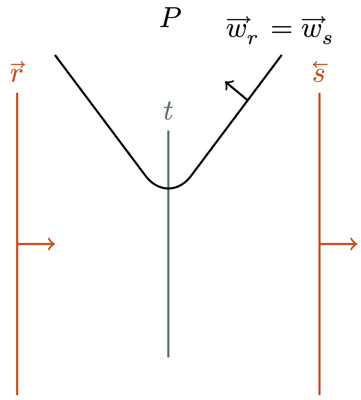

If $\vec{w}_r = \vec{w}_s$ crosses t, then, by the choice of t, that neither $\vec{w}_r \wedge \vec{t}$ nor $\vec{w}_r \wedge \overleftarrow{t}$ has order less than $|t|$, thus $\vec{w}_r \vee \vec{t}$ and $\vec{w}_r \vee \overleftarrow{t}$ both have order at most the order of $\vec{w}_r$. By the strong robustness of P applied to $\vec{w}_r, \overleftarrow{w}_r \wedge \overleftarrow{t}$ and $\overleftarrow{w}_r \wedge \vec{t}$, we know that either $\vec{w}_r \vee \vec{t} \in P$ or $\vec{w}_r \vee \overleftarrow{t} \in P$. However, both contradict the $\leqslant$-maximal choice of $\vec{w}_r$. So, instead $\vec{w}_r$ is nested with t, that is, t has an orientation $\vec{t}$ such that $\vec{t} \leqslant \vec{w}_r$, so t is weakly oriented by P, as claimed. □

Note that the assumption that our profiles are *strongly* robust is essential in this argument, for example for the case $\vec{w}_r = \vec{w}_s$: If we only assume robustness, we can not conclude that P contains either $\vec{w}_r \vee \vec{t}$ or $\vec{w}_r \vee \overleftarrow{t}$ and thus would not obtain a contradiction.

The next step is to verify that shifting with a splice as in Lemma 6.25 maps stars in $\mathcal{F}_d$ to stars in $\mathcal{F}_d$, which will prove that U is critically $\mathcal{F}_d$-seperable:

Lemma 6.26. *Let r and s_0 be separations which are weakly oriented by every profile in $\mathcal{P}$ and suppose that $\vec{s}_0$ us a splice for $\vec{r}$. Let $\sigma \in \mathcal{F}_d$ be a star which contains a separation $\vec{x} \geqslant \vec{r}$. Then the shift $\sigma^{\vec{s}_0}_{\vec{x}}$ of σ from $\vec{x}$ to $\vec{s}_0$ is again an element of $\mathcal{F}_d$.*

Proof. Since $\vec{s}_0$ is a splice for $\vec{r}$, by Lemma 6.8, $\vec{s} \vee \vec{s}_0$ has at most the order of $\vec{s}$ for every $\vec{s} \geqslant \vec{r}$.

Let σ be any star in $\mathcal{F}_d$ containing a separation $\vec{x} \geqslant \vec{r}$. By the above, if $\sigma \subseteq \vec{S}_k$ for some k, then also the shift $\sigma^{\vec{s}_0}_{\vec{x}}$ is a subset of $\vec{S}_k$. Hence by Lemma 6.12, every profile in U which contains $\sigma^{\vec{s}_0}_{\vec{x}}$ also contains σ. Now if some profile P contains σ, then P orients every separation in $\sigma_{\vec{x}}^{\vec{s}_0}$, and thus either P contains the inverse of some separation in $\sigma_{\vec{x}}^{\vec{s}_0}$ or $\sigma_{\vec{x}}^{\vec{s}_0} \subseteq P$.

Hence, by Lemma 6.17 it is enough to show that every profile from $\mathcal{P}$ which, for some $\vec{y} \in \sigma$, weakly contains $\overleftarrow{y}$ also weakly contains $\overleftarrow{y}'$ for some separation $\vec{y}' \in \sigma_{\vec{x}}^{\vec{s}_0}$.

So suppose such a profile P, for some $\vec{y} \in \sigma$, weakly contains $\overleftarrow{y}$ and suppose that this is witnessed by $\overleftarrow{w}_y \in P$. If $\vec{r} \leqslant \vec{y}$, then $\overleftarrow{y}$ is shifted onto $\overleftarrow{y} \wedge \overleftarrow{s}_0$ and therefore $\overrightarrow{w}_y$ also witnesses that P weakly contains $\overleftarrow{y} \wedge \overleftarrow{s}_0$ while $\vec{y} \vee \vec{s}_0 \in \sigma_{\vec{x}}^{\vec{s}_0}$. Thus we may suppose that $\vec{r} \leqslant \overleftarrow{y}$ and therefore that $\overleftarrow{y}$ is shifted onto $\overleftarrow{y} \vee \vec{s}_0$.

If P weakly orients s_0 as $\overleftarrow{s}_0$, then P also weakly contains $\vec{y} \wedge \overleftarrow{s}_0 \leqslant \overleftarrow{s}_0$ while $\overleftarrow{y} \vee \vec{s}_0 \in \sigma_{\vec{x}}^{\vec{s}_0}$.

Thus we may suppose that P weakly orients s_0 as $\vec{s}_0$, witnessed by $\overrightarrow{w}_0 \in P$.

By our assumptions on s_0 we know that the order of $\vec{s}_0 \wedge \overleftarrow{w}_y$ is at least the order of s_0 and thus, by submodularity, $\overleftarrow{s}_0 \wedge \overrightarrow{w}_y$ has order at most the order of w_y, i.e., it is oriented by P. By Lemma 6.24 applied to $\overrightarrow{w}_0, \overleftarrow{w}_y \in P$ and $\overleftarrow{s}_0 \wedge \overrightarrow{w}_y$ we can therefore conclude that P contains $\vec{s}_0 \vee \overleftarrow{w}_y$, i.e., P weakly contains $\overleftarrow{y} \vee \vec{s}_0 \leqslant \vec{s}_0 \vee \overleftarrow{w}_y$. $\square$

In order to use our stronger tangle–tree duality theorem, Theorem 6.23, with our set $\mathcal{F}_d$ of stars to obtain a tree of tangles for strongly robust profiles it only remains for us to show that this application cannot result in an $\mathcal{F}_d$-tangle. We do so in the following two lemmas.

Lemma 6.27. *For every profile P in $\vec{U}$ and every set $\mathcal{P}'$ of strongly robust profiles in $\vec{U}$ distinguishable from P, there exists a nested set N which distinguishes P efficiently from all the profiles in $\mathcal{P}'$.*

Proof. For every profile $Q \in \mathcal{P}'$ pick a $\leqslant$-minimal separation $\vec{s}_Q \in P$ which efficiently distinguishes Q from P. We claim that the set N consisting of all these separations s_Q is nested and therefore as claimed.

So suppose that this is not the case, so $\vec{s}_Q$ and $\vec{s}_{Q'}$, say, cross. We may assume without loss of generality that $|\vec{s}_Q| \leqslant |\vec{s}_{Q'}|$. Now $\vec{s}_Q \vee \vec{s}_{Q'}$ has order at least the order of $\vec{s}_{Q'}$ since otherwise, by the profile property, $\vec{s}_Q \vee \vec{s}_{Q'}$ would also distinguish P and Q' and would thus contradict the fact that $\vec{s}_{Q'}$ did so efficiently. Thus $|\vec{s}_Q \wedge \vec{s}_{Q'}| \leqslant |\vec{s}_Q|$.

Now Q' orients $\vec{s}_Q$ and it cannot contain $\overleftarrow{s}_Q$ since then, by the profile property, $\vec{s}_Q \wedge \vec{s}_{Q'}$ would also distinguish P and Q' efficiently and would therefore contradict the $\leqslant$-minimal choice of $\vec{s}_{Q'}$.

Thus $\vec{s}_Q \in Q'$. Now $|\overleftarrow{s}_Q \wedge \vec{s}_{Q'}| > |\vec{s}_{Q'}|$ since otherwise, again by the profile property, $\overleftarrow{s}_Q \wedge \vec{s}_{Q'}$ contradicts the $\leqslant$-minimal choice of $\vec{s}_{Q'}$.

Thus, by submodularity, $|\vec{s}_Q \wedge \overleftarrow{s}_{Q'}| < |\vec{s}_Q|$ and $|\vec{s}_Q \wedge \vec{s}_{Q'}| \leqslant |\vec{s}_Q|$. But, by strong robustness, either $\overleftarrow{s}_Q \vee \vec{s}_{Q'}$ or $\overleftarrow{s}_Q \vee \overleftarrow{s}_{Q'}$ is in Q. In particular, $\vec{s}_Q \wedge \overleftarrow{s}_{Q'}$ or $\vec{s}_Q \wedge \vec{s}_{Q'}$ efficiently distinguishes P and Q and therefore contradicts the $\leqslant$-minimal choice of $\vec{s}_Q$. $\square$

Unlike for structurally submodular separation systems in Lemma 6.14 or efficient distinguishers in Lemma 6.21, in this setup we can not necessarily find a star in $\mathcal{F}_d$ which is contained in O but in no profile in $\mathcal{P}$ for every orientation O of U which not include any profile in our set $\mathcal{P}$ of strongly robust profiles. This is because we require that every profile in $\mathcal{P}$ weakly orients a separation in our star outwards, but the stars constructed in Lemma 6.21, for example, do not necessarily have this property. Thus we are going to, instead, find a star σ contained in both O and exactly one profile from $\mathcal{P}$. Since each such star also lies in $\mathcal{F}_d$, this will be enough to ensure that our application of Theorem 6.23 does not result in an $\mathcal{F}_d$-tangle.

Lemma 6.28. *For every consistent orientation O of $\vec{U}$ and every set $\mathcal{P} \neq \emptyset$ of distinguishable strongly robust profiles in $\vec{U}$ there exists a star σ in $\mathcal{F}_d$ contained in O.*

Proof. Pick a star σ (not necessarily from $\mathcal{F}_d$) with the following properties:

(i) $\sigma \subseteq O$.

(ii) σ is contained in at least one profile in $\mathcal{P}$.

(iii) Property $\mathsf{Eff}(P)$ is satisfied for every profile $P \in \mathcal{P}$ such that $\sigma \subseteq P$.

(iv) Every $P \in \mathcal{P}$ either contains σ or weakly contains $\overleftarrow{s}$ for some separation $\vec{s} \in \sigma$.

(v) For every separation $\vec{s} \in \sigma$ and any profile $\sigma \subseteq P$ there exists a profile $Q \in \mathcal{P}$ such that s is a efficient P–Q-distinguisher.

Note that the empty set is such a star. Let us further assume that we choose our star σ fulfilling (i)–(v) so that as few profiles in $\mathcal{P}$ as possible contain σ.

If only one profile contains σ, then $\sigma \in \mathcal{F}_d$ is as desired, so let us suppose for a contradiction that there are at least two such profiles.

Pick two such profiles $P_1, P_2 \supseteq \sigma$ such that the order of an efficient P_1–P_2-distinguisher is as small as possible. Pick an efficient P_1–P_2-distinguisher s which crosses as few elements of σ as possible. O orients s, say $\vec{s} \in O$. If s is nested with σ, the maximal elements of $\sigma \cup \{\vec{s}\}$ form a star violating the definition of σ: Every profile containing this new star also contains σ. To see that (iii) is fulfilled, note that there is no profile $P \supseteq \sigma$ in $\mathcal{P}$ such that $\vec{s} \in P$ for which there is a $\vec{s}'$ of lower order than $\vec{s}$ such that $\vec{s} \leqslant \vec{s}' \in P$, since such an $\vec{s}'$ would be a distinguisher of lower order than $\vec{s}$ for some pair of profiles containing σ, contrary to the coice of s.

Thus we may assume that s is not nested with σ, say s crosses $\vec{t} \in \sigma$. Since, by (v), there is some profile $Q \ni \overleftarrow{t}$ for which t is an efficient P_1–Q-distinguisher, we know that at least one of $\vec{s} \wedge \vec{t}$ and $\overleftarrow{s} \wedge \vec{t}$ has order at least the order of t: Otherwise this would contradict the fact that t is an efficient P_1–Q-distinguisher by robustness (if $|t| < |s|$) or the profile property (if $|s| \leqslant |t|$) of Q.

Thus by submodularity the order of at least one of $\vec{s} \vee \vec{t}$ and $\overleftarrow{s} \vee \vec{t}$ is at most the order of s and that separation is therefore also an efficient P_1–P_2-distinguisher (by the profile property and consistency), which would make it a better choice for s, a contradiction.

Thus σ contains precisely one profile and therefore, by construction, $\sigma \in \mathcal{F}_d$. $\square$

Together with Theorem 6.23, these lemmas give a proof of a tree-of-tangles theorem for strongly robust profiles of different orders in a submodular universe. This theorem does not give efficient distinguishers; we will deal with efficiency in a later step.

Theorem 6.29 (Tree-of-tangles theorem for different orders)**.** *Let $(\vec{U}, \leqslant, ^*, \vee, \wedge, |\cdot|)$ be a submodular distributive universe of separations. Then for every distinguishable set $\mathcal{P}$ of strongly robust 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 distinguished by some separation in T;

(ii) for any profile $P \in \mathcal{P}$, any maximal $\vec{s} \in P \cap \vec{T}$ and any $\vec{s}\,' \in P$ such that $\vec{s} \leqslant \vec{s}\,'$ we have $|\vec{s}| \leqslant |\vec{s}\,'|$.

Proof. By Lemma 6.25 and Lemma 6.26 the set U is critically $\mathcal{F}_d$-separable for the set $\mathcal{F}_d$ defined above. Thus we can apply Theorem 6.23. This can, by Lemma 6.28, not result in an $\mathcal{F}_d$-tangle, thus there is an U-tree over $\mathcal{F}_d$. By Lemma 6.4 we may assume this U-tree to be irredundant. The set of separations associated to edges of this tree is then a nested set T.

Every profile in $\mathcal{P}$ induces a consistent orientation of T, since all the separations in T are weakly oriented by every profile in $\mathcal{P}$. The maximal elements of this orientation form a star σ_P in $\mathcal{F}_d$, and this star is a subset of P by the definition of $\mathcal{F}_d$.

To see that T distinguishes every pair of profiles in $\mathcal{P}$, consider two profiles P and Q in $\mathcal{P}$. These two profiles cannot induce the same orientation of T, since then $\sigma_P = \sigma_Q$ would be a subset of both P and Q, contradicting the definition of $\mathcal{F}_d$. Thus some $\vec{s} \in \sigma_P$ witnesses that P weakly orients some $\vec{t} \in \sigma_Q$ as $\overleftarrow{t}$ and, vice versa, $\vec{t}$ witnesses that Q weakly contains $\overleftarrow{s}$. Of these two separations s and t, the one of lower order is thus a P–Q-distinguisher in T.

Property (ii) is then immediate from the definition of $\mathcal{F}_d$. $\square$

Note that the nested set constructed in Theorem 6.29 does not yet necessarily distinguish any two profiles *efficiently*. However, we can use Theorem 6.16 in combination with Theorem 6.29 to obtain such a set:

Theorem 6.30 (Efficient tree-of-tangles theorem for different order profiles)**.** *Let $(\vec{U}, \leqslant, ^*, \vee, \wedge, |\cdot|)$ be a submodular distributive universe of separations. Then for every distinguishable set $\mathcal{P}$ of strongly robust profiles in $\vec{U}$ there is a nested set $T = T(\mathcal{P}) \subseteq U$ of separations such that every two profiles in $\mathcal{P}$ are efficiently distinguished by some separation in T.*

Proof. Let k be the maximal order of a profile in U. Let T be the U-tree over $\mathcal{F}_d$ from the proof of Theorem 6.29. We consider the $\subseteq$-maximal subtrees T_i of T with the property that no internal node of T_i corresponds to a profile in $\mathcal{P}$. Clearly $T = \bigcup_{i=1}^{m} T_i$ and no two T_i share an edge.

We are going to simultaneously replace each of the nested sets of separations corresponding to the T_is with other separations in such a way that the resulting set of separations is still nested and we ensured that every pair of profiles contained in some T_i is efficiently distinguished by this new set of separations.

So, given some T_i, let $\mathcal{P}_i$ be the set of profiles in $\mathcal{P}$ living, in T, in one of the leaves of T_i. Let $\vec{L}_i$ be the set of all separations associated to one of the directed edges adjacent and pointing away from such a leaf. Note that $\vec{L}_i$ is a star. For every $\vec{s} \in \vec{L}_i$ let $P_s \in \mathcal{P}_i$ be the unique profile corresponding to a leaf of T_i and containing $\overleftarrow{s}$.

It is easy to check that for any two profiles P and Q in $\mathcal{P}_i$ there is a efficient P–Q-distinguisher t which is nested with all of $\vec{L}_i$: Pick one t which is nested with as many separations from $\vec{L}_i$ as possible. Now t cannot cross an $\vec{s} \in \vec{L}_i$ such that $P_s = P$ or $P_s = Q$, as in that case, for $\vec{t} \in P_s$, either $\vec{t} \vee \overleftarrow{s}$ or $\vec{t} \wedge \overleftarrow{s}$ would, by submodularity, consistency and the profile property, be an efficient P-Q--distinguisher and as such contradict the choice of t by Lemma 2.1. If on the other hand t crosses some $\vec{s} \in \vec{L}_i$, such that $P_s \notin \{ P, Q \}$, then not both of $\overleftarrow{s} \vee \vec{t}$ and $\overleftarrow{s} \vee \overleftarrow{t}$ can have order less than the order of s by the profile property since, by property (ii), there is no $\overleftarrow{s}' \in P_s$ such that $\overleftarrow{s} \leqslant \overleftarrow{s}'$ and $|\vec{s}| > |\vec{s}'|$. Thus the order of either $\vec{s} \vee \vec{t}$ or $\vec{s} \vee \overleftarrow{t}$ is at most the order of t, however by Lemma 6.24 and the fish lemma, Lemma 2.1, this separation then contradicts the choice of t.

Moreover, there exists such an efficient P–Q-distinguisher t which has an orientation $\vec{t}$ such that $\vec{s} \leqslant \vec{t}$ for every $\vec{s} \in \vec{L}_i$: Otherwise $\overleftarrow{s} \leqslant \vec{t}$ for some orientation of t and if neither $P = P_s$ nor $Q = P_s$, then both P and Q would weakly orient t as $\overleftarrow{t}$ since they weakly contain $\vec{s}$. On the other hand if $P = P_s$, say, then, again by property (ii), the order of t is at least the order of s, thus s itself would be the required efficient P–Q-distinguisher.

Now consider, for every T_i, the set U^i of all separations t in U nested with $\vec{L}_i$ and fulfilling the additional property of having, for every $\vec{s} \in \vec{L}_i$, an orientation such that $\vec{s} \leqslant \vec{t}$, i.e., U^i is the set of all separations in U inside of $\vec{L}_i$. $\vec{U}^i$ is closed under $\vee$ and $\wedge$ in $\vec{U}$ by the fish Lemma 2.1, thus the restriction of U to U^i is again a submodular universe of separations.

Given any $\vec{s} \in \vec{L}_i$, the down-closure of $\overleftarrow{s}$ is a regular profile of U^i. Note that every efficient distinguisher for the profiles induced by $\overleftarrow{s}_1$ and $\overleftarrow{s}_2 \in \overleftarrow{L}_i$ on U^i is also an efficient distinguisher of P_{s_1} and P_{s_2}.

By Theorem 6.16 applied to the set of all separations of order less than k in U^i, we thus find a U^i-tree $\hat{T}^i$ over $\mathcal{F}_e$ (defined for $\mathcal{P}_i$). The corresponding nested set N_i efficiently distinguishes all these profiles induced by some $\overleftarrow{s}_i \in \overleftarrow{L}_i$.

But now the nested N given by $\bigcup_{i=1}^{m}(N_i \cup L_i)$ is as desired: It is easy to see that this set is nested and every N_i efficiently distinguishes any two profiles in $\mathcal{P}_i$. Moreover, we only ever changed separations inside of $\vec{L_i}$ for every T_i.

The set N also contains an efficient P–Q-distinguisher for profiles P and Q in different T_is: A profile R whose node in T lies on the path between the nodes containing P and Q, respectively, also does so in the tree induced by N. Thus, if we have efficient distinguishers for P and R and for R and Q, respectively, in N, then one of the two is also an efficient P–Q-distinguisher. An inductive application of this argument proves the claim, that the set N efficiently distinguishes any two profiles in $\mathcal{P}$. $\square$

7. Tangle-finding and a tree-of-tangles algorithm

This chapter presents a practical approach to tangle computation. It is loosely based on the preprint note [35], which is joint work with Christian Elbracht and Jakob Kneip; this chapter contains new, unpublished variations of those ideas as indicated throughout the text and in Appendix A.4.

If we want to apply tangles to identify dense clusters in real-world data, then we will quickly run into the issue that the time and space requirements for precise computation of most kinds of tangles on large ground sets can be immense. While good work has been put into making the computation as efficient as possible, such as the algorithm by Grohe and Schweitzer for finding the tangles and a tree of tangles for a connectivity function [51], there is only so much that can be done, since deciding whether a graph has tree-width k, where both k and the graph are input variables, is NP-complete [1] and this decision problem is linked to the existence of tangles in the graph: the existence of a high-order tangle implies high tree-width and high tree-width is witnessed by a high-order tangles, up to a constant factor. [64]

Instead of calculating those hard-to-compute tangles of a connectivity function, in this section we will utilize the generality that abstract separation systems have to offer: we will work not with *the* tangles in some submodular universe $\vec{U}$, but with just some small separation system $\vec{S}$ which we think of as some 'representative sample' of $\vec{U}$. There is a variety of approaches how such a set $\vec{S}$ might be obtained for any given set-up, e.g. from a rough[1] MinCut-approximation algorithm, see [33] for a demonstration of just some approaches.

We will first introduce the most obvious[2] algorithm for finding $\mathcal{F}$-tangles of some separation system in Section 7.2. This algorithm assumes that we are given the separation system $\vec{S}$ as an enumeration where the separations are in order of increasing value of the order function; this algorithm will help us build an understanding for the setting. We then, in Section 7.3, present a more refined tangle-search algorithm to which the separations can be provided in arbitrary order, e.g. in the order that they come out of our separation sampling process. This

[1]In practice it seems to be the case that you want *locally* minimal cuts, in some sense. A too good MinCut-algorithm would just give you the global minimal cut which, alone, is not helpful for tangles.

[2]That is, apart from independently checking for every orientation whether it is a tangle.

algorithm can run in parallel to our sampling process, enabling us to keep track of the tangle structure uncovered by the so-far sampled separations while avoiding the redundant computation that would come from re-running the algorithm from Section 7.2 for every new separation. This also opens the option of having the tangle structure inform the sampling of new separations.

In the worst case both of these algorithms have a runtime which is exponential in $|\vec{S}|$. As an extreme example, if $\mathcal{F} = \emptyset$, then every one of the $2^{\vec{S}}$ subsets of $\vec{S}$ avoids $\mathcal{F}$ and thus the number of tangles is exponential in $|\vec{S}|$, requiring the runtime of the algorithm to be exponential as well. Practically, the runtime is determined mainly by the number of tangles which exist, and any sensible data analysis using tangles should adjust $\mathcal{F}$ so that the number of tangles is not too large in order to get meaningful information out of tangles anyway. Moreover, the number of separations that we consider is a parameter that we have control over, so we can make trade-offs between runtime and amount of detail that we recover.

Finally, in Section 7.4, we show how one can build trees of tangles in the separation-sampling approach. A randomly sampled separation system $\vec{S}$ will be unlikely to allow for a nested set of separations that distinguishes all its tangles, but careful expansion of $\vec{S}$ can make that possible.

7.1. Setting and Terminology

In this chapter we consider a separation system $\vec{S}$ together with an order function $|\cdot| : \vec{S} \to \mathbb{R}_{\geqslant 0}$. Usually $\vec{S}$ will be a 'sampled' subset of some submodular universe $\vec{U}$, which is also where the order function comes from. (However, this will only become a proper prerequisite for Section 7.4.) Analogous to tangles *in* a submodular universe, we will define tangles *in* S.

Let $\mathcal{F}$ be some fixed set of sets of separations. For $r \in \mathbb{R} \cup \infty$, let us denote by $S_{<r}$ the set $\{\, s \in S : |s| < r \,\}$. Note that $S_{<r}$ is defined as a subset of S and not of some universe which S might be a subset of, thus we avoid the usual S_k-notation here. An *r-tangle in S* shall be an $\mathcal{F}$-avoiding orientation τ of $S_{<r}$.[3] A *tangle in S* is an r-tangle in S for any r, and the *order* of a tangle τ in S is the maximum $r \in \mathbb{R} \cup \infty$ for which τ is an r-tangle. A *maximal tangle* in S is one which is not a subset of any higher-order tangle in S.

For the presentation of the algorithms, we will always assume that the order function $|\cdot|$ takes distinct values on distinct unoriented separations. This is to avoid some rather technical bookkeeping, which would get in the way of understanding the main parts of the algorithms. However, all the algorithms presented can be augmented with a certain amount of bookkeeping to be able to do away with

[3]We do not demand consistency separately here. Instead, if consistency is required, we encode this by letting $\mathcal{F}$ contain all inconsistent pairs of separations. This is a minor detail which will simplify the upcoming presentation of our algorithms.

this assumption. From a theoretical standpoint we can also apply Corollary 9.9 to perturb our order function into one which takes distinct values on distinct unoriented separations, as we did for unrabelling order-induced submodular separation systems in Section 9.2. Practically, you would first have to find an efficient way of calculating a tie-breaker like the function from Corollary 9.9 for any given setup.

Our algorithms will be given access to $\mathcal{F}$ not as a complete list but rather as a membership oracle, i.e., a procedure which decides for any given subset of $\vec{S}$ whether it is an element of $\mathcal{F}$ or not. A full list of the elements of $\mathcal{F}$ would be far to long to be practical in most cases. Thus, such an oracle access is preferable and aligns well with how tangle conditions are usually defined. In most settings we know that the elements of $\mathcal{F}$ have a bounded size of at most some $m \in \mathbb{N}$.[4] We will mark the places in the algorithms where this knowledge can be used to avoid some unnecessary computation.

7.2. The tangle search tree

The first algorithm – an algorithm for finding all the tangles in S – can already be found in our note on generic tangle algorithms [35] and was applied in [33] as joint work with machine learning scientists. It is also a very straightforward algorithm which naturally arises from the definitions; a pseudo code representation can be seen in Algorithm 1.

To find out whether a given orientation τ of $\vec{S}$ avoids $\mathcal{F}$ we have to check all the subsets of τ for membership in $\mathcal{F}$. (At least those of size $\leqslant m$.) Since any subset of τ also avoids $\mathcal{F}$, we can also formulate this recursively: if $\tau = \tau' \cup \{ \vec{s} \}$ then τ avoids $\mathcal{F}$ if, and only if, τ' avoids $\mathcal{F}$ and no $<m$-element subset of τ together with $\vec{s}$ is in $\mathcal{F}$.

Now if we want to find all tangles of S, we can use this recursive approach bottom-up to solve the problem of finding tangles by a binary tree search, as follows. Let $S = \{ s_1, \dots, s_n \}$, enumerated with increasing order, and denote $S^i := \{ s_1, \dots, s_i \}$. Then we can define a binary tree T on the vertex set $\bigsqcup_{0 \leqslant i \leqslant n} 2^{\vec{S}^i}$, where every $\tau \in 2^{\vec{S}^i}$ (for $i > 0$) has the parent $\tau^- := \tau \cap \vec{S}^{i-1}$. If τ is a tangle of $\vec{S}^i$ then τ^- is one of $\vec{s}^{i-1}$. Thus, the tangles in S form a binary subtree of T, and we can perform a tree search (depth-first or breadth-first) of this binary subtree to find all the tangles of $\vec{S}$, starting at the trivial tangle $\emptyset$.[5]

In this way, we do not only find all the tangles of S but all tangles of all the $\vec{S}^i$s. Thus, since we chose the enumeration of $\vec{S}$ to be non-decreasing with respect to the order function $|\cdot|$, this gives us not only all tangles *of* S, but all tangles *in* S. This computation then yields the whole *hierarchy of tangles* in $\vec{S}$ represented by

[4]For typical choices of $\mathcal{F}$, like in the cases of profiles and of tangles in graphs, we have $m = 3$.

[5]Unless $\emptyset \in \mathcal{F}$ which would be silly, since then no $\mathcal{F}$-tangles can possibly exist.

TangleSearch($[s_1, \dots, s_n], \mathcal{F}$)

$\mathcal{T}_{\mathsf{max}} \leftarrow \emptyset$
$\mathcal{T}_0 \leftarrow \{\emptyset\}$
for $i = 1 \dots n$
    $\mathcal{T}_i \leftarrow \emptyset$
    for each $\tau \in \mathcal{T}_{i-1}$
        isMaximal $\longleftarrow$ True
        for each orientation $\vec{s}$ of s_i
            if $\forall X \subseteq \tau\colon \{\vec{s}\} \cup X \notin \mathcal{F}$ (*Checking* X *of size* $\leqslant m-1$ *suffices.*)
                $\mathcal{T}_i \leftarrow \mathcal{T}_i \cup \{\tau \cup \{\vec{s}\}\}$
                isMaximal $\longleftarrow$ False
        if *isMaximal*
            $\mathcal{T}_{\mathsf{max}} \leftarrow \mathcal{T}_{\mathsf{max}} \cup \{\tau\}$
$\mathcal{T}_{\mathsf{max}} \leftarrow \mathcal{T}_{\mathsf{max}} \cup T$
return $\mathcal{T}_{\mathsf{max}}$

Algorithm 1: Pseudo-code for a depth-first tangle search yielding the set of maximal tangles.

the binary tree. The leaves of this tree are the maximal tangles in S and the whole tree can easily be reconstructed from just them.

7.3. On-line tangle search

Suppose we are not given access to all of $\vec{S}$ at once but that it is given to us one separation at a time in an uncontrollable order. Maybe these separations come from a randomized process, and it is unreasonable to demand that separations are given to us in increasing order. We now want to be able to compute the tangle hierarchy *on-line*, i.e., going through the separations in the order they are given to us and maintaining a list of the tangles found so far, in the given subset of $\vec{S}$. This will also come in handy later, in the context of the tree of tangles, where we will 'discover' separations which had not been considered during tangle search.

We encode the order in which the separations are given to us as some fixed enumeration $S = \{ s_1, s_2, \dots, s_n \}$ which need not respect $|\cdot|$, and we shall write S^i for $\{ s_1, \dots, s_i \}$.

Our aim is to find the maximal tangles in S inductively: given the maximal tangles of S^i and their respective orders, we consider s_{i+1} and compute the maximal tangles of S^{i+1}.

$\textsc{UpdateTangle}(\tau, k, \vec{s}, \mathcal{F})$

$[\vec{s}_1, \dots, \vec{s}_\ell] \leftarrow \tau$ (*Enumeration of τ with increasing order.*)
$i_0 \leftarrow \mathsf{min}(\{\, i : |\vec{s}| < |\vec{s}_i| \,\} \cup \{\, \ell + 1 \,\})$ (*Determine insertion point for s.*)
for each $X \subseteq \{\, \vec{s}_1, \dots, \vec{s}_{i_0-1} \,\}$ (*Checking X of size $\leqslant m-1$ suffices.*)
  if $X \cup \{\, \vec{s} \,\} \in \mathcal{F}$ ($\{\, \vec{s}_1, \dots, \vec{s}_{i_0-1}, \vec{s} \,\}$ *includes a forbidden set.*)
    return $[\vec{s}_1, \dots, \vec{s}_{i_0-1}]$, $|s|$
for $i = i_0, \dots, n$
  for each $X \subseteq \{\, \vec{s}_1, \dots, \vec{s}_{i-1} \,\}$ (*Checking X of size $\leqslant m-2$ suffices.*)
    if $X \cup \{\, \vec{s}_i, \vec{s} \,\} \in \mathcal{F}$ ($\{\vec{s}_1, \dots, \vec{s}_{i_0-1}, \vec{s}, \vec{s}_{i_0}, \dots \vec{s}_i\}$ *includes a forbidden set.*)
      return $[\vec{s}_1, \dots, \vec{s}_{i_0-1}, \vec{s}, \vec{s}_{i_0}, \dots \vec{s}_{i-1}]$, $|s_i|$
return $[\vec{s}_1, \dots, \vec{s}_{i_0-1}, \vec{s}, \vec{s}_{i_0}, \dots \vec{s}_i]$, k

$\textsc{UpdateTangles}(\mathcal{T}, s, \mathcal{F})$

$\mathcal{T}' \leftarrow \emptyset$
for each $(\tau, k) \in \mathcal{T}$
  $\tau_1, k_1 \leftarrow \textsc{UpdateTangle}(\tau, k, \vec{s}, \mathcal{F})$
  $\tau_2, k_2 \leftarrow \textsc{UpdateTangle}(\tau, k, \overleftarrow{s}, \mathcal{F})$
  $\mathcal{T}' \leftarrow \mathcal{T} \cup \{\, (\tau_1, k_1), (\tau_2, k_2) \,\}$
return $\textsc{MaximalElements}(\mathcal{T})$

$\textsc{TangleSearchOnLine}(S, \mathcal{F})$

$\mathcal{T} \leftarrow \{\, (\emptyset, \infty) \,\}$
for each $s \in S$
  $\mathcal{T} \leftarrow \textsc{UpdateTangles}(\mathcal{T}, s, \mathcal{F})$
return $\mathcal{T}$

Algorithm 2: The procedures that make up the on-line tangle search algorithm. The procedure $\textsc{MaximalElements}$ is trivial and not listed; it determines the maximal elements with respect to $(k, \tau) \leqslant (k', \tau') :\Leftrightarrow \tau \subseteq \tau'$.

A pseudo code version of this algorithm is displayed in Algorithm 2. In the $(i+1)$th step of the algorithm, for every maximal tangle τ in S^i and every orientation $\vec{s}$ of s_{i+1}, we do the following: let r be the order of τ. Set $\tau' = \tau \cup \{\vec{s}\}$ and determine the largest r' for which $\tau'_{r'} := \tau' \cap \vec{S}^{i+1}_{<r'} = \{\vec{s} \in \tau' : |s| < r'\}$ is an r'-tangle, i.e., the largest r' for which the elements of τ' of order $< r'$ contain no subset in $\mathcal{F}$. This r' is at least $|s|$ and at most r. In-between the two it may take any value which is the order of a separation in S^i. We can perform these checks for membership in $\mathcal{F}$ in-order, so that we do not test sets for membership in $\mathcal{F}$ unnecessarily, see Algorithm 2.

Finally, the maximal elements among all such τ' are the maximal tangles in S^{i+1}, as desired.

7.4. Building a tree of tangles with partial information

Once we have found the tangles in some $\vec{S}$, we might want to use them to split up our structure into a tree of tangles. Suppose we are working in a submodular universe $\vec{U}$ and the set $\mathcal{F}$ includes all those two- and three-element subsets of $\vec{U}$ which would violate consistency, the profile property (P), or the robustness property (R). If the separation system that we ran the algorithm on was some $\vec{S} = \vec{S}_k \subseteq \vec{U}$, then the profiles in $\vec{S}$ are robust profiles in $\vec{U}$ and our simplifying assumption that the order function takes distinct values on different separations makes it very easy to find a tree of tangles: recall that for such a set of robust profiles there exists a nested set of separations which *efficiently* distinguishes all the profiles by Theorem 5.8. Now,, in this set-up, every pair of profiles is efficiently distinguished precisely by the unique separation of the smallest order among the separation on which the profiles differ. So, if we pick for every pair of profiles this unique separation, we will automatically obtain a nested set.

What can we do if we are not given $\vec{S}_k$ as a whole, but only a fragment of it? Closing our separation system $\vec{S}$ under all possible corners in $\vec{U}$ can be prohibitively expensive in terms of time and memory, let alone running a tangle search on it. But, using our insights from the splinter lemma in Section 5.2, we can still build a tree of tangles where we extend $\vec{S}$ by only those separations which are necessary to achieve a nested set, as we will now explain.

First, we take for every pair of tangles in $\vec{S}$ an efficient distinguishing separations in $\vec{S}$, as above; these might not form nested set. We will now make a corner-taking argument along the lines of the arguments we made for the applications of the splinter lemma in Section 5.3, as follows.

Say there is a pair s, t of separations that cross, s distinguishing some tangles $\tau_{\vec{s}}, \tau_{\overleftarrow{s}}$ efficiently and t distinguishing $\tau_{\vec{t}}, \tau_{\overleftarrow{t}}$ efficiently, where each contains the separation indicated by its index. Assume w.l.o.g. that $|s| < |t|$, then $\tau_{\vec{t}}$ and $\tau_{\overleftarrow{t}}$ both orient s, and that, by the efficiency of t, both orient s the same, as $\vec{s}$ say. By the submodularity of $|\cdot|$, either one of $\vec{s} \vee \vec{t}$ and $\vec{s} \vee \overleftarrow{t}$ has a lower order than t, or both $\overleftarrow{s} \vee \vec{t}$ and $\overleftarrow{s} \vee \overleftarrow{t}$ have a lower order than s. In either case, the fact that $\mathcal{F}$ enforces robust profiles implies that one of those lower-order corners has not been considered during our tangle search, i.e., one of the tangles does not orient it but orients some separation of higher order. We thus run the UpdateTangles procedure of Algorithm 2 with that separation and start anew with building a distinguishing set. We repeat this process until the set of efficient distinguishers comes out nested.

This process terminates; at the latest once we have run UpdateTangles with all the separations in U. Take note that UpdateTangles can change the number of tangles both ways: a tangle might split in two, if both orientations of the newly considered separation are possible; and a tangle might disappear by being truncated in UpdateTangle so that it comes out as a subset of another tangle and is then eliminated by MaximalElements.

The result of this process is a separation system $\vec{S}' \subseteq \vec{U}$, which is a superset of the separation system $\vec{S}$ that we started with, the set $\mathcal{T}'$ of all $\mathcal{F}$-tangles in $\vec{S}'$ and a nested set of separations $N \subseteq S'$ which distinguishes all the tangles in $\mathcal{T}'$ efficiently.

It is possible to adapt this algorithm to allow distinct separations of the same order. If we allow multiple separations of the same order, then we have to choose between multiple possible efficient distinguisher-separations for each pair of tangles. Even if $\vec{S}$ was some $\vec{S}_k$, not every possible such choice would be nested. We can however follow our proof of Lemma 5.5 to make them nested, resolving one crossing pair at a time. In general, our approach requires not discarding the distinguishing set and starting anew after every UpdateTangles, but to keep track of one choice of distinguishers and making changes to it only as necessary. One can then follow the proof of Lemma 5.5, to gradually make the distinguishing set 'more and more nested'. See our note [35] for a detailed, and very technical, variant of this approach which, additionally, goes out of it's way to update as few tangles as necessary to build a nested set – at the cost of the tangles becoming tangles of $\subseteq$-incomparable subsystems $\vec{S} \subseteq \vec{U}$.

Part III.

Submodularity

8. Submodularity in separation systems

In this chapter we analyse and differentiate three key notions of submodularity. The contents of this chapter are based on the preprint [37] which is joint work together with Christian Elbracht and Jakob Kneip. Section 8.6 is my own addition.

A central property that is required of separation systems in almost every context is some form of *submodularity*: a property needed to make the separation system 'rich enough' to prove the desired theorems, for example the tree-of-tangles theorem [26]. Our aim for this part of the thesis is to study, and relate, the various forms of submodularity.

Originally, in [26, 28, 47, 64], it was not only required that $\vec{S}$ was part of a universe of separations but also that there exist a submodular order function $f\colon \vec{U} \to \mathbb{R}_{\geqslant 0}$ and $k \in \mathbb{R}_{\geqslant 0}$ such that

$$\vec{S} = \vec{S}_k := \{ \vec{s} \in \vec{U} : f(\vec{s}) < k \}.$$

Diestel, Erde, and Weißauer [25] showed that the theorems of tangle theory could also be deduced without relying on such an order function, demanding instead just one structural property of $\vec{S} \subseteq \vec{U}$ which in the case of the sets $\vec{S}_k$ is imposed by the submodularity of f: that for all $\vec{r}, \vec{s} \in \vec{S}$ at least one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ is also in $\vec{S}$. Note that this structural property of $\vec{S}$ is measured externally: in the universe $\vec{U}$, where the join $\vec{r} \vee \vec{s}$ and the meet $\vec{r} \wedge \vec{s}$ are taken. To reflect this, we say that $\vec{S} \subseteq \vec{U}$ is *submodular in* $\vec{U}$. Whenever a submodular order function f on $\vec{U}$ and a number k exist such that $\vec{S} = \vec{S}_k$ for this order function, we say that the submodularity of $\vec{S}$ in $\vec{U}$ is *order-induced* in $\vec{U}$.

Much of the work presented in this thesis, particularly in Part II, relies heavily on such structural submodularity of separation systems, rather than on the existence of a submodular order function. Indeed, separation systems which are submodular in some universe of separations form the most relevant class of separation systems nowadays, and the most general theorems of abstract tangle theory are formulated in their context [25, 26, 28, 34].

The most natural structural notion of submodularity, however, is simply to call a separation system $\vec{S}$ *submodular* if any two separations $\vec{r}, \vec{s} \in \vec{S}$ have either a supremum $\vec{r} \vee \vec{s}$ or an infimum $\vec{r} \wedge \vec{s}$ in $\vec{S}$. Unlike in our earlier definition of submodularity for $\vec{S}$ in some universe $\vec{U}$, the question now is whether such infima and suprema exist – not whether they lie in $\vec{S}$.

Note that every separation system $\vec{S}$ that is submodular in some universe $\vec{U}$ of separations is also submodular also in this latter sense, since every infimum or supremum of $\vec{r}, \vec{s} \in \vec{S}$ in $\vec{U}$ which (by submodularity in $\vec{U}$) also lies in $\vec{S}$, is also the supremum or infimum of $\vec{r}$ and $\vec{s}$ in $\vec{S}$. Submodularity of a separation system $\vec{S}$, as defined locally in $\vec{S}$ itself, is therefore a weakening of submodularity in some surrounding universe of separations.

One can then ask whether this weaker kind of submodularity still suffices as a basis for the theorems of tangle theory, which traditionally assume that the separation system $\vec{S}$ whose tangles are studied is submodular in some universe $\vec{U}$. Our first result, which we prove in Section 8.2, shows that it does:

⇝ p. 173 **Theorem 29.** *Every submodular separation system is submodular in some universe of separations.*

Theorem 29 allows us to apply the main theorems of tangle theory to separation systems which are known to be submodular only in the weaker local sense, without the need to re-prove them under this weaker assumption.

In Sections 8.3 and 8.4 we turn our attention to the question of when the submodularity of a separation system in a universe $\vec{U}$ is always induced by a submodular order function on $\vec{U}$. In Section 8.3 we prove that it need not be:

⇝ p. 177 **Theorem 30.** *There exists a separation system $\vec{S}$ which is submodular in a universe $\vec{U}$ of set bipartitions, but whose submodularity in $\vec{U}$ is not induced by a submodular order function on $\vec{U}$.*

More precisely, we present a necessary condition for the submodularity of a separation system in a universe $\vec{U}$ to be order-induced in $\vec{U}$, and use this to give concrete examples of systems which are submodular in some universe $\vec{U}$ of separations but whose submodularity is not order-induced in this $\vec{U}$.

In Section 8.4 we consider another aspect of order-induced submodularity. Whether the submodularity in a universe $\vec{U}$ of a separation system is order-induced or not depends, a priori, on the choice of $\vec{U}$. As a simple example, consider the case that a separation system $\vec{S}$ is submodular in a universe $\vec{U}$ of separations, and that $\vec{U}$ is a subuniverse of some larger universe $\vec{U}'$ of separations. Then $\vec{S}$ is submodular also in $\vec{U}'$. If the submodularity of $\vec{S}$ in $\vec{U}$ is witnessed by some submodular order function on $\vec{U}$, we may ask whether we can extend this function to $\vec{U}'$ to witness that $\vec{S}$ is submodular also in $\vec{U}'$. We show that this can be done in some cases. The general question of whether it is always possible to extend such a witnessing submodular order function to a larger universe remains open.

Finally, in Section 8.5, we present two decomposition theorems for separation systems that are submodular in distributive universes. Our first decomposition theorem allows us to write every such separation system $\vec{S}$ as a (not necessarily disjoint) union of three smaller ones, each of which is not only again submodular in the same universe, but is also closed under taking existing corners in $\vec{S}$. Thus, we cover $\vec{S}$ by smaller, simpler, 'spanned' subsystems. To prove this, we introduce a variation of Birkhoff's representation theorem for universes of separations instead of lattices. Moreover, in our decomposition theorem, the subsystems can be chosen disjoint, unless the separation system to be decomposed is one of set bipartitions.

Separation systems that are submodular in the universe $\vec{U}$ of bipartitions of a set V cannot be decomposed disjointly into submodular subsystems. Indeed, every non-empty subsystem would have to contain the separations $(V, \emptyset)$ and $(\emptyset, V)$, since these form opposite corners of every pair of inverse separations. By submodularity in $\vec{U}$, one of these – and hence also the other, as its inverse – would have to lie in this subsystem.

Separation systems of set bipartitions are, however, very concrete and better understood than the more general abstract separation systems. We may view these bipartition systems as the 'elementary building blocks' which make up the separation systems that are submodular in distributive universes. Applying our decomposition theorem repeatedly, for as long as disjoint decompositions are possible, we can thus break down every separation system that is submodular in a distributive universe into those elementary subsystems.

Theorem 35. *Every separation system $\vec{S}$ which is submodular in some distributive universe $\vec{U}$ of separations is a disjoint union of corner-closed subsystems $\vec{S}_1, \dots, \vec{S}_n$ of $\vec{S}$ (which are thus also submodular in $\vec{U}$) each of which can be corner-faithfully embedded into a universe of bipartitions.* ⇝ p. 187

Specifically, these subsystems are the equivalence classes of the relation $\sim$ on $\vec{S}$ where $\vec{s} \sim \vec{t}$ if and only if $\vec{s} \wedge \overleftarrow{s} = \vec{t} \wedge \overleftarrow{t}$ in $\vec{U}$.

In the chapter after this one, Chapter 9, we will introduce the unravelling problem, which is concerned with yet another, very particular, property in which the three kinds of submodularity differ.

8.1. Lattice terminology and tools

For this chapter we will use some additional terminology and insights from lattice theory wherein we follow the definitions of the textbook of Davey and Priestley [16]. We also introduce generalizations of the submodularity concepts from separation systems and universes of separations to posets and lattices.

Recall, that a lattice is *distributive* if it satisfies the distributive laws, that is for all $a, b, c \in L$ we have that $a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c)$ and $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$. A typical example of a distributive lattice is the *subset lattice* of a set V, which consists of the subsets of V ordered by $\subseteq$. Here the join of two sets is their union, the meet is their intersection.

In fact, all finite distributive lattices can be represented as a set of subsets where $\vee$ and $\wedge$ coincide with union and intersection. This is a fundamental result of lattice theory known as the *Birkhoff representation theorem*, which we can state after the following additional definitions: a non-bottom element $x \in L$ is *join-irreducible* if whenever $x = a \vee b$ for some $a, b \in L$, then $x \in \{ a, b \}$. The set of all join-irreducible elements of L is denoted $\mathcal{J}(L)$ and forms a partially ordered set with the order inherited from L. Given a partially ordered set $(P, \leqslant)$, the down-closed sets in P form a distributive lattice with $\subseteq$ as the partial order, union as join and intersection as meet. This lattice is denoted as $\mathcal{O}(P)$.

Theorem 8.1 (Birkhoff representation theorem; cf. [16, 5.12])**.** *Let L be a finite distributive lattice. The map $\eta\colon L \to \mathcal{O}(\mathcal{J}(L))$ defined by $\eta(a) = \{ x \in \mathcal{J}(L) : x \leqslant a \} = \lceil a \rceil_{\mathcal{J}(L)}$ is an isomorphism of lattices.*

Given a lattice L, any subset $P \subseteq L$ together with the restrictions of $\vee$ and $\wedge$ (as *partial functions*) is called a *partial lattice*, cf. [52].

In analogy to submodularity of separation systems in universes, we now introduce submodular subsets of lattices. Let L be a lattice and $P \subseteq L$ a subset of L. We say that P is *submodular in L* if for all $a, b \in P$ at least one of $a \vee b$ and $a \wedge b$ is in P as well.

If $f\colon L \to \mathbb{R}_{\geqslant 0}$ is a submodular function on L, then the sets $f^{-1}([0, x))$ are submodular in L. If for some submodular P in L there exists some submodular function f and some x such that $P = f^{-1}([0, x))$, then we say that the submodularity of P in L is *induced by f (and x)*, or that the submodularity of P in L is *order-induced.*[1]

We say that a poset P (in and of itself) is *submodular* if any two elements $a, b \in P$ have a supremum in P or an infimum in P.

8.2. Witnessing submodularity externally

The traditional, external, notions of submodularity always require our separation system $\vec{S}$ to be part of a universe $\vec{U}$ of separations, even though, often, we are interested in $\vec{S}$ and substructures therein only and do not particularly care about the shape of $\vec{U}$. The only reason for keeping this ambient universe $\vec{U}$ around is that we need to be able to express joins and meets of elements of $\vec{S}$, and decide whether these lie in- or outside of $\vec{S}$. The mathematical arguments exploiting the submodularity of $\vec{S}$ never truly make use of $\vec{U}$, but only of the knowledge that at least one of two opposing corner separations, $\vec{r} \vee \vec{s}$ and $\overleftarrow{r} \vee \overleftarrow{s} = (\vec{r} \wedge \vec{s})^*$, is always present in $\vec{S}$, for all $\vec{r}, \vec{s} \in \vec{S}$.

As discussed above, the simplest, most general, and thus most natural, form of submodularity for a separation system is intrinsic from its poset structure, where a poset $P = (P, \leqslant)$ is submodular if all pairs $a, b \in P$ have a supremum or an infimum

[1] This term is used to emulate the notation for separation system. The 'order'-part of order-induced submodularity, i.e., the symmetry of the submodular function is not applicable to lattices.

in P. Yet almost all theorems in the theory of abstract separation systems are phrased in terms of some form of submodularity which is external, in some universe of separations, even when that universe bears no particular relevance on the result.

In this section we offer a way out: a method by which the submodularity of some $\vec{S}$ in itself can be reflected into a suitable universe of separations in such a way, that its submodularity is expressed externally. If $\vec{S}$ is a separation system which is submodular in some universe $\vec{U}$, then $\vec{S}$ is also submodular on its own. Here, we will show a converse to this: if a separation system $\vec{S}$ is submodular on its own, then we can construct a universe $\vec{U}$ which contains an isomorphic copy of the separation system $\vec{S}$, i.e., there is an *embedding* of $\vec{S}$ into $\vec{U}$. We can choose this embedding in such a way, that the pre-existing joins and meets inside $\vec{S}$ are preserved. More precisely, if $\vec{r}$ and $\vec{s}$ have a supremum $\vec{t}$ in $\vec{S}$, then after embedding $\vec{S}$ into $\vec{U}$ we will have $\vec{t} = \vec{r} \vee \vec{s}$, where the latter is measured in $\vec{U}$. Thus our $\vec{U}$ is chosen so that $\vec{S}$ is *submodular in* $\vec{U}$.

Theorem 8.2. *For every separation system $\vec{S}$, finite or infinite, there exists a universe $\vec{U}$ of separations and an embedding $\varphi\colon \vec{S} \to \vec{U}$, with the property that $\varphi(\vec{t}) = \varphi(\vec{r}) \vee \varphi(\vec{s})$ if and only if $\vec{t}$ is the supremum of $\vec{r}$ and $\vec{s}$ in $\vec{S}$, and likewise $\varphi(\vec{u}) = \varphi(\vec{r}) \wedge \varphi(\vec{s})$ if and only if $\vec{u}$ is the infimum of $\vec{r}$ and $\vec{s}$ in $\vec{S}$. Moreover, if $\vec{S}$ is finite, then $\vec{U}$ can be chosen to be finite.*

In particular, if $\vec{S}$ is submodular, then $\varphi(\vec{S})$ is submodular in $\vec{U}$.

The heavy lifting of Theorem 8.2's proof is done by employing the Dedekind–MacNeille-completion [62], a lattice theoretic tool with which one can embed an arbitrary poset into a suitable lattice while preserving any pre-existing finite joins and meets.[2] Our task then is to equip the resulting completion of the poset $\vec{S}$ with an involution which turns it into a universe of separations, and which makes the embedding of $\vec{S}$ into its Dedekind–MacNeille-completion an isomorphism onto its image.

To define this Dedekind–MacNeille-completion, we follow the notation of [16]. Let P be any poset, finite or infinite. Given a subset $X \subseteq P$ we write X^ℓ for the set of lower bounds of X in P: the set of all $p \in P$ such that $p \leqslant x$ for all $x \in X$. Similarly we write X^u for the set of all upper bounds of X in P. To improve readability we will omit braces when concatenating these operations, e.g., we shall write $X^{u\ell}$ rather than $(X^u)^\ell$.

The *Dedekind–MacNeille-completion* of P is then given by

$$\mathbf{DM}(P) := \{\, X \subseteq P : X^{u\ell} = X \,\}$$

using $\subseteq$ as the partial order. A result by MacNeille [62] asserts that $\mathbf{DM}(P)$ is indeed a lattice and, moreover, the map $\varphi\colon P \to \mathbf{DM}(P)$ given by

$$\varphi(p) := \{p\}^\ell$$

[2] The Dedekind–MacNeille-completion is more commonly used for infinite lattices, where it is used to embed a lattice into a complete lattice, hence the name. It is also known as the *completion by cuts.*

is an embedding of the poset P into $\mathbf{DM}(P)$ with the property that $\varphi(r)$ is the supremum (resp. infimum) of $\varphi(p)$ and $\varphi(q)$ if and only if r is the supremum (resp. infimum) of p and q in P. (Compare [16, Theorem 7.40].)

To build some intuition about the Dedekind–MacNeille-completion, observe that for a singleton $\{p\}$, the set $\{p\}^u$ is simply the up-closure $\lfloor p \rfloor$ of p in P. Moreover an element q of P is a lower bound of the up-closure of some p precisely if $q \leqslant p$, and hence $\{p\}^{u\ell} = \{p\}^\ell = \lceil p \rceil$. In particular, when applying any series of ^{u} and ${}^\ell$ to a singleton set $\{p\}$, only the very last operation is relevant: for instance $\{p\}^{\ell u \ell} = \{p\}^\ell$, which shows that the map φ indeed takes its image in $\mathbf{DM}(P)$.

Let us now prove Theorem 8.2.

Proof of Theorem 8.2. Let $\vec{S} = (\vec{S}, \leqslant, {}^*)$ be a separation system. Let $\vec{U} = \mathbf{DM}(\vec{S})$ be the Dedekind–MacNeille-completion of the poset $\vec{S}$ with the embedding $\varphi \colon \vec{S} \to \vec{U}$ given by $\varphi(\vec{s}) = \{\vec{s}\}^\ell$.

For a set $X \subseteq \vec{S}$ we write X^* for the point-wise involution $\{\, \overleftarrow{x} : \vec{x} \in X \,\}$ of X. For readability we shall extend our convention to omit braces to include *, ^{u}, and ${}^\ell$. Clearly $X^{**} = X$ for all $X \subseteq \vec{S}$.

We define an involution ${}^\circledast$ on $\vec{U}$ by letting

$$X^\circledast := X^{u*}$$

and claim that this turns $\vec{U}$ into a universe of separations and φ into an isomorphism of separation systems between $\vec{S}$ and its image in $\vec{U}$. To verify this claim we need to ascertain the following: that ${}^\circledast$ takes its image in $\vec{U} = \mathbf{DM}(\vec{S})$, that ${}^\circledast$ is an involution, that ${}^\circledast$ is order-reversing, and finally that φ commutes with the involution, i.e., that $\varphi(\vec{s})^\circledast = \varphi(\overleftarrow{s})$.

Before we do this, observe that since the involution * of $\vec{S}$ is order-reversing we have

$$X^{u*} = X^{*\ell} \qquad \text{and} \qquad X^{\ell *} = X^{*u}$$

for all $X \subseteq \vec{S}$. We shall be using these two equalities throughout the remainder of the proof.

To see that ${}^\circledast$ takes its image in $\vec{U}$, note that for $X \in \vec{U}$ we have

$$(X^\circledast)^{u\ell} = X^{u*u\ell} = X^{u\ell*\ell} = X^{*\ell} = X^{u*} = X^\circledast ,$$

where the third equality used the definition of $\vec{U} = \mathbf{DM}(\vec{S})$ to infer $X^{u\ell} = X$. Thus we indeed have $X^\circledast \in \vec{U}$ by definition of $\vec{U} = \mathbf{DM}(\vec{S})$.

The map ${}^\circledast$ is an involution since

$$(X^\circledast)^\circledast = X^{u*u*} = X^{u\ell**} = X^{u\ell} = X ,$$

for $X \in \vec{U}$, using again the definition of $\vec{U} = \mathbf{DM}(\vec{S})$.

To see that $^{\circledast}$ is order-reversing let $X, Y \in \vec{U}$ with $X \subseteq Y$ be given; we need to show that $X^{\circledast} \supseteq Y^{\circledast}$. From $X \subseteq Y$ it follows that $X^u \supseteq Y^u$, which in turn implies $X^{u*} \supseteq Y^{u*}$. Thus indeed $X^{\circledast} \supseteq Y^{\circledast}$.

We now show that $\varphi(\vec{s})^{\circledast} = \varphi(\overleftarrow{s})$ for all $\vec{s} \in \vec{S}$. So let $\vec{s} \in \vec{S}$ be given. Recall that $\{\vec{s}\}^{\ell u} = \{\vec{s}\}^u$ and $\varphi(\vec{s}) = \{\vec{s}\}^{\ell}$. Using this equality we find that

$$\varphi(\vec{s})^{\circledast} = \varphi(\vec{s})^{u*} = \{\vec{s}\}^{\ell u*} = \{\vec{s}\}^{u*} = \{\vec{s}\}^{*\ell} = \{\overleftarrow{s}\}^{\ell} = \varphi(\overleftarrow{s})\,,$$

as claimed.

Since φ preserves the existing pairwise suprema and infima of the poset $\vec{S}$, it is thus the desired embedding. $\square$

We can phrase Theorem 8.2 more concisely, as follows:

Theorem 29. *Every submodular separation system is submodular in some universe of separations.*

8.3. Structural submodularity which is not order-induced

In this section we deal with the question whether the submodularity of a submodular subsystem $\vec{S} \subseteq \vec{U}$ of a universe $\vec{U}$ is always induced by some submodular order function f on $\vec{U}$, i.e., that $\vec{S} = \vec{S}_k$ for some k. We will answer this question in the negative, even for distributive $\vec{U}$, and thus show that submodularity in a universe is a proper generalization of order-induced submodularity.

We consider the question first for partial lattices $P \subseteq L$ which are submodular in some lattice L. Recall that these are partial lattices $P \subseteq L$ such that for any two points $a, b \in P$ at least one of $a \vee b$ and $a \wedge b$ (taken in L) is in P.

One way to show that the submodularity of a given partial lattice is not order-induced is to find a sequence $a_1, a_2, \dots, a_n$ of elements of a lattice L so that every submodular function f on L for which P is an S_k would need to satisfy $f(a_1) < f(a_2) < \cdots < f(a_n) < f(a_1)$. Such a sequence may be found by finding a directed cycle in a digraph D on L where we draw an edge from a to b whenever every suitable submodular function on L needs to satisfy $f(a) > f(b)$.

This motivates the following definition: for $P \subseteq L$ we define the *dependency digraph* $D = (L, E)$ of P as a directed graph where (a, b) is an edge in E if and only if one of the following holds:

- $a \in L \setminus P$ and $b \in P$;
- $a, b \in P$ and there is some $c \in P$ such that either
  - $b = a \vee c$ and $a \wedge c \notin P$, or
  - $b = a \wedge c$ and $a \vee c \notin P$;
- $a, b \notin P$ and there is some $c \in P$ such that either

- $b = a \vee c$ and $a \wedge c \notin P$, or
- $b = a \wedge c$ and $a \vee c \notin P$.

Let us first show that given an order-induced submodular partial lattice $P \subseteq L$, the edges in the dependency digraph indeed witness that their start vertex has higher order than their end vertex.

Lemma 8.3. *If $P \subseteq L$ is order-induced submodular, witnessed by some f and k, and (a, b) is an edge in the dependency digraph of P, then $f(a) > f(b)$.*

Proof. Let (a, b) be an edge in the dependency digraph. If $a \in L \setminus P$ and $b \in P$ then $f(a) > f(b)$ since f induces the submodularity of P in L.

If $a, b \in P$ we may assume without loss of generality that the edge between a and b exists because of some $c \in P$ with $b = a \vee c$ and $a \wedge c \notin P$.

Because f induces the submodularity of P in L we have $f(a \wedge c) > f(c)$. Since f is submodular

$$f(a \vee c) + f(a \wedge c) \leqslant f(a) + f(c),$$

and hence $f(b) = f(a \vee c) < f(a)$, as required.

Similarly, if $a, b \notin P$ we may assume without loss of generality that the edge between a and b exists because of some $c \in P$ with $b = a \vee c$ and $a \wedge c \notin P$.

Because f induces the submodularity of P in L we have $f(a \wedge c) > f(c)$. Again, since f is submodular

$$f(a \vee c) + f(a \wedge c) \leqslant f(a) + f(c),$$

and hence $f(b) = f(a \vee c) < f(a)$, as required. □

Thus a directed cycle in the dependency digraph is an obstruction to the order-induced submodularity of P.

Corollary 8.4. *If the dependency digraph of P contains a directed cycle then P is not order-induced submodular.*

Since every cycle in the dependency digraph D of P is completely contained in either $D[P]$ or $D[L \setminus P]$, we sometimes consider these two subgraphs independently of each other, naming them the *inner dependency digraph* $D[P]$ and the *outer dependency digraph* $D[L \setminus P]$.

Each cycle in the dependency digraph has length at least 3:

Lemma 8.5. *Let $P \subseteq L$ be submodular in L, then the dependency digraph of P contains no directed cycle of length 2.*

Proof. As stated above, a cycle of length 2 cannot contain one vertex in P and one in $L \setminus P$. Thus if the dependency digraph D contains a cycle of length 2 between a and b, then by the definition of the dependency digraph a and b are comparable in

$\leqslant$, so $a \leqslant b$, say. Note that either $a, b \in P$ or $a, b \notin P$. In either case, as (a, b) is an edge in D, there exists a $c \in P$ such that $a \vee c = b$ and $a \wedge c \notin P$. Similarly, there exists a $d \in P$ such that $b \wedge d = a$ and $b \vee d \notin P$.

If $c \leqslant d$ then $d \geqslant a$ and $d \geqslant c$ and thus $a \vee c = b \leqslant d$ contradicting the assumption that $b \vee d \notin P$. Similarly, if $d \leqslant c$ then $d \leqslant c \leqslant b$, again contradicting the assumption. Hence, c and d are incomparable and thus $c \vee d \in P$ or $c \wedge d \in P$, as $c, d \in P$ and P is submodular in L. However, $b = a \vee c \leqslant d \vee c$, thus $d \vee c \geqslant b$, hence $d \vee c \geqslant b \vee d$, but also $d \vee c \leqslant d \vee b$ as $c \leqslant b$, and thus $d \vee c = d \vee b \notin P$. And similarly, $a = d \wedge b \geqslant d \wedge c$, thus $d \wedge c \leqslant a \wedge c$ but also $d \wedge c \geqslant a \wedge c$ and thus $d \wedge c = a \wedge c \notin P$.

Thus, D cannot contain a cycle of length 2. $\square$

Using the dependency digraph, we can give an example of a lattice L together with a partial lattice $P \subseteq L$ which is submodular in L, but where this submodularity is not order-induced. Our example will use a universe of separations as its lattice, and a submodular separation system for the partial lattice.

In fact, our example consists of oriented bipartitions (equivalently: subsets) on a set of six elements. The Hasse diagram of this example is displayed in Fig. 8.1; a formal description follows.

Consider the universe $\vec{U} = \vec{\mathcal{B}}(V)$ of bipartitions of $V = \{\, a, b, c, d, e, f \,\}$. In there we consider the separation system $\vec{S}$ consisting of the orientations of the following unoriented bipartitions:

$$\begin{aligned} S = \{\, & \{\, \emptyset, V \,\}, \\ & \{\, \{b\}, \{a, c, d, e, f\} \,\}, \{\, \{d\}, \{a, b, c, e, f\} \,\}, \{\, \{f\}, \{a, b, c, d, e\} \,\}, \\ & \{\, \{a, b\}, \{c, d, e, f\} \,\}, \{\, \{c, d\}, \{a, b, e, f\} \,\}, \{\, \{e, f\}, \{a, b, c, d\} \,\}, \\ & \{\, \{a, b, c\}, \{d, e, f\} \,\}, \{\, \{a, b, f\}, \{c, d, e\} \,\}, \{\, \{a, e, f\}, \{b, c, d\} \,\} \,\}. \end{aligned}$$

It is easy to see that $\vec{S}$ is submodular in $\vec{U}$. However, the dependency digraph of $\vec{S}$ in $\vec{U}$ contains the directed cycle

$$\begin{aligned} (\{a, b, c, d\}, \{e, f\}) &\to (\{a, b\}, \{c, d, e, f\}) \to (\{a, b, e, f\}, \{c, d\}) \to (\{e, f\}, \{a, b, c, d\}) \\ &\to (\{c, d, e, f\}, \{a, b\}) \to (\{c, d\}, \{a, b, e, f\}) \to (\{a, b, c, d\}, \{e, f\}). \end{aligned}$$

For example, there is an arc between $(\{a, b, c, d\}, \{e, f\})$ and $(\{a, b\}, \{c, d, e, f\})$ since

$$(\{a, b, c, d\}, \{e, f\}) \wedge (\{a, b, f\}, \{c, d, e\}) = (\{a, b\}, \{c, d, e, f\})$$

and

$$(\{a, b, c, d\}, \{e, f\}) \vee (\{a, b, f\}, \{c, d, e\}) = (\{a, b, c, d, f\}, \{e\}),$$

but $(\{a, b, c, d, f\}, \{e\})$ is not an element of $\vec{S}$. The existence of the remaining arcs in the cycle can be checked similarly.

This example proves the following theorem:

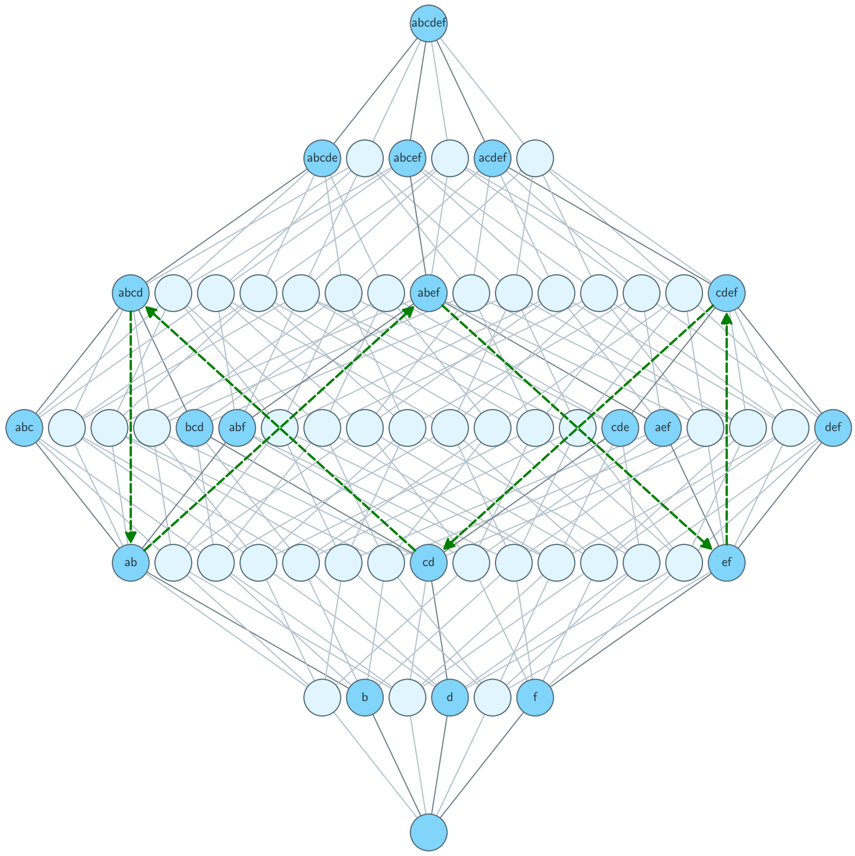

Figure 8.1.: The Hasse diagram of $\vec{U}$ from Theorem 30. For readability, only points in $\vec{S}$ are labelled and only one side of each bipartition is denoted.

Theorem 30. *There exists a separation system $\vec{S}$ which is submodular in a universe $\vec{U}$ of set bipartitions, but whose submodularity in $\vec{U}$ is not induced by a submodular order function on $\vec{U}$.*

One might wonder if every example of a partial lattice with a cycle in its dependency digraph actually contains a cycle in the *inner* dependency digraph. This is not the case, as an example we show the Hasse digram of such a lattice in Fig. 8.2 and indicate the partial lattice inside this lattice as well as the cycle in the dependency digraph.

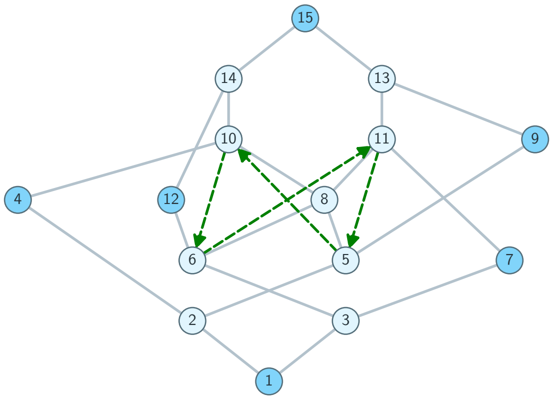

Figure 8.2.: The dark blue elements form a partial lattice, which does not contain a cycle in the inner dependency digraph, however the green dashed edges form a cycle in the outer dependency digraph

However, we are not aware of any examples of submodular separation systems whose submodularity in a universe is not order-induced and whose dependency digraph is acyclic:

Question 8.6. Does there exists a separation system $\vec{S} \subseteq \vec{U}$ which is submodular in $\vec{U}$, such that the dependency digraph of $\vec{S}$ does not contain a cycle, but the submodularity of $\vec{S}$ in $\vec{U}$ is not order-induced?

We can ask the same question for a submodular partial lattice:

Question 8.7. Does there exists a partial lattice $P \subseteq L$ which is submodular in the lattice L such that the dependency digraph of P does not contain a cycle, but the submodularity of P in L is not order-induced?

These two questions are, in fact, equivalent. To see this, observe that a positive answer to Question 8.6 implies a positive answer to Question 8.7: if there exists a separation system $\vec{S} \subseteq \vec{U}$ which is submodular in $\vec{U}$, such that the dependency digraph of $\vec{S}$ does not contain a cycle, but the submodularity is not order-induced, then we can consider $\vec{S}$ as a partial lattice inside the lattice $\vec{U}$ which still does not contain a cycle in its dependency digraph. However, if $g\colon \vec{U} \to \mathbb{R}_{\geqslant 0}$ would be a submodular function witnessing that $\vec{S}$ is order-induced submodular as a partial lattice with some $k \in \mathbb{R}_{\geqslant 0}$, then we could consider the function f given by $f(\vec{s}) = g(\vec{s}) + g(\overleftarrow{s})$ for every $\vec{s} \in \vec{S}$, which would then be a submodular order function for $\vec{U}$, and f and $2k$ induce the submodularity of $\vec{S}$ in the universe $\vec{U}$.

On the other hand, if there exists a partial lattice $P \subseteq L$ which is submodular in the lattice L such that the dependency digraph of P does not contain a cycle, but the submodularity is not order-induced, we can construct a universe $\vec{U}$ and a submodular subsystem $\vec{S} \subseteq \vec{U}$, so that the dependency digraph of $\vec{S}$ does not contain a cycle, but the submodularity of $\vec{S}$ in $\vec{U}$ is not order-induced, as follows: let L' be a copy of L with reversed partial order (i.e., the poset-dual of L). We let $\vec{U}$ be the disjoint union $L \sqcup L'$, where we additionally declare $\vec{r} \leqslant \vec{s}$ for all $\vec{r} \in L$ and $\vec{s} \in L'$. The involution on $\vec{U}$ is defined by mapping an element of L to its respective copy in L' and vice versa. It is easy to see that this is a universe of separations and that $\vec{S} = P \cup P'$ (where $P \subseteq L$ is as above and $P' \subseteq L'$ is the image of P in L') is a submodular subsystem of $\vec{U}$. [3] Moreover, $\vec{S}$ is not order-induced submodular, since we can restrict any witnessing submodular order function on $\vec{U}$ to a submodular function on L, which would then witness that the submodularity of P in L is order-induced.

The dependency digraph of $\vec{S}$ cannot contain a cycle either, since any such cycle would result in a cycle in the dependency digraph of L or L': every edge in the dependency digraph of $\vec{U}$ either is also an edge in the dependency digraph of L or L', or is an edge between L and L' which needs to be an edge between an element of $\vec{U} \setminus \vec{S}$ and $\vec{S}$. Thus, given any cycle in the dependency digraph of $\vec{U}$ which meets both L and L', we can consider a maximal subpath of this cycle contained in L; there then needs to be a directed edge in the dependency of L between the last and the first vertex of this path.

8.4. Extending a submodular function

Our aim in this section is to better understand for what kind of submodular separation systems the submodularity is order-induced. We investigate inhowfar the existence of a submodular function depends on the surrounding universe $\vec{U}$, that is, if we have an order function f which induces the submodularity of some $\vec{S}$ in a subuniverse $\vec{U'} \subseteq \vec{U}$, we ask whether we can extend f to $\vec{U}$ in such a way that it induces the submodularity of $\vec{S}$ in $\vec{U}$.

We give partial answers to this question: firstly that submodular functions can be extended in this way from an *interval* in a universe and, secondly, that for every

[3] Note, that in $\vec{U}$ every separation is either small or co-small, i.e., for every $\vec{s} \in \vec{U}$ either $\vec{s} \leqslant \overleftarrow{s}$ or $\overleftarrow{s} \leqslant \vec{s}$.

subuniverse $\vec{U}'$ of a universe $\vec{U}$ there exists a submodular function f and some k, such that $\vec{U}' = \vec{S}_{f,k} = f^{-1}([0,k))$, that is $\vec{U}'$ is order-induced submodular in $\vec{U}$.

It suffices to first consider these problems for submodular functions on lattices, rather than submodular order functions on universes of separations: if $f' : \vec{U}' \to \mathbb{R}_{\geqslant 0}$ is a submodular function on $\vec{U}' \supseteq \vec{U}$ which agrees on $\vec{U}$ with some submodular *order* function $f\colon \vec{U} \to \mathbb{R}_{\geqslant 0}$, then we can define a submodular *order* function $\bar{f}$ on $\vec{U}'$ which agrees with f by setting

$$\bar{f}(\vec{s}) := \frac{f'(\vec{s}) + f'(\overleftarrow{s})}{2}.$$

We will then easily see that, in both cases, this function is as desired.

For the first theorem, recall that an *interval* in a lattice L is, for some $x, y \in L$, a subset $[x,y] = \{ s \in L : x \leqslant s \leqslant y \}$. Every such interval forms a sublattice. The following result shows that we can extend a submodular function defined on an interval.

Theorem 8.8. *Let L be a lattice and $L' = [x,y] \subseteq L$ an interval in L. Suppose that $f\colon L' \to \mathbb{R}_{\geqslant 0}$ is a submodular function on L' with maximum value k. Then there exists a submodular function $g\colon L \to \mathbb{R}_{\geqslant 0}$ such that $g(z) = f(z)$ for all $z \in L'$ and $g(z) > k$ for all $z \notin L'$.*

Proof. Let us denote as $L^{\downarrow}$ the set of all $z \in L \setminus L'$ such that $z \leqslant y$, as $L^{\uparrow}$ the set of all $z \in L \setminus L'$ such that $z \geqslant x$ and as $L^{\leftrightarrow}$ the set of all $z \in L \setminus L'$ such that neither $z \leqslant y$ nor $z \geqslant x$. Note that $L^{\downarrow}, L^{\uparrow}, L^{\leftrightarrow}$ and L' together form a partition of L.

For $z \in L$ such that $z \leqslant y$ we define its *down-level* $\mathsf{dl}(z)$ recursively as follows: assign $\mathsf{dl}(\bot) = 0$ for the bottom element $\bot$ of L and set $\mathsf{dl}(z) := \max\{ \mathsf{dl}(z') + 1 : z' < z \}$ for all other $z \in L$. Similarly, for $z \in L$ such that $z \geqslant x$ we define its *up-level* $\mathsf{ul}(z)$ recursively: we assign $\mathsf{ul}(\top) = 0$ for the top element $\top$ of L and set $\mathsf{ul}(z) := \max\{ \mathsf{ul}(z') + 1 : z' > z \}$ for all other $z \in L$.

Let ℓ be the maximum possible level (up or down) and let $M = 2^{\ell} \cdot k > k$. We now define g as follows:

$$g(z) = \begin{cases} f(z) & \text{if } z \in L', \\ M \cdot (2 - 2^{-\mathsf{dl}(z)}) & \text{if } z \in L^{\downarrow}, \\ M \cdot (2 - 2^{-\mathsf{ul}(z)}) & \text{if } z \in L^{\uparrow}, \\ 4 \cdot M & \text{if } z \in L^{\leftrightarrow}. \end{cases}$$

To verify that this function is submodular we distinguish the possible cases which can occur for two incomparable elements $a, b \in L$. Note that in the case of comparable elements, submodularity is trivially satisfied, so we suppose they are incomparable.

The case $a, b \in L^{\leftrightarrow}$.

By construction, the maximal value of g is $4 \cdot M$, thus

$$g(a \vee b) + g(a \wedge b) \leqslant 4 \cdot M + 4 \cdot M = g(a) + g(b).$$

The case $a \in L^{\uparrow}, b \in L^{\leftrightarrow}$.
By the definition of $L^{\uparrow}$, we have $a \vee b \in L^{\uparrow}$ and $\mathsf{ul}(a) > \mathsf{ul}(a \vee b)$, thus

$$\begin{aligned} g(a \vee b) + g(a \wedge b) &\leqslant M \cdot (2 - 2^{-\mathsf{ul}(a \vee b)}) + 4 \cdot M \\ &< M \cdot (2 - 2^{-\mathsf{ul}(a)}) + 4 \cdot M = g(a) + g(b). \end{aligned}$$

The case $a \in L^{\downarrow}, b \in L^{\leftrightarrow}$.
Analogous to the above.

The case $a \in L', b \in L^{\leftrightarrow}$.
By the definition of $L^{\uparrow}$, we have, since $a \vee b \geqslant a \geqslant x$, that $a \vee b \in L^{\uparrow} \cup L'$ and similarly, $a \wedge b \in L^{\downarrow} \cup L'$. Thus, we have

$$g(a \vee b) + g(a \wedge b) \leqslant 2M + 2M \leqslant g(b) \leqslant g(a) + g(b).$$

The case $a, b \in L^{\uparrow}$.
Suppose without loss of generality that $\mathsf{ul}(a) \leqslant \mathsf{ul}(b)$. By the definition of $L^{\uparrow}$ and ul, we have $a \vee b \in L^{\uparrow}$ and $\mathsf{ul}(a \vee b) < \mathsf{ul}(a)$. Furthermore $a \wedge b \in L^{\uparrow} \cup L'$, so in any case $g(a \wedge b) < 2M$. We calculate

$$\begin{aligned} g(a \vee b) + g(a \wedge b) &< M \cdot (2 - 2^{-(\mathsf{ul}(a)-1)}) + 2M \\ &= 4M - M(2^{-\mathsf{ul}(a)} + 2^{-\mathsf{ul}(a)}) \\ &\leqslant 4M - M(2^{-\mathsf{ul}(a)} + 2^{-\mathsf{ul}(b)}) = g(a) + g(b). \end{aligned}$$

The case $a, b \in L^{\downarrow}$.
Analogous to the above.

The case $a \in L^{\downarrow}, b \in L^{\uparrow}$.
By construction $a \wedge b \in L^{\downarrow}$ and $a \vee b \in L^{\uparrow}$. Moreover, by the definition of g we have $g(a \wedge b) \leqslant g(a)$ and $g(a \vee b) \leqslant g(b)$ and thus

$$g(a \wedge b) + g(a \vee b) \leqslant g(a) + g(b).$$

The case $a \in L', b \in L^{\uparrow}$.
By the definition of $L^{\uparrow}$, we have $a \vee b \in L^{\uparrow}$. Moreover $ul(a \vee b) < ul(b)$, by the definition of g and choice of M, we thus have $g(a \vee b) \leqslant g(b) - k$. Additionally, $g(a \wedge b) \in L'$, since $x \leqslant a \wedge b$ and $a \wedge b \leqslant a \leqslant y$. Thus, by the definition of k, we have $g(a \wedge b) \leqslant g(a) + k$ and thus

$$g(a \vee b) + g(a \wedge b) \leqslant g(b) - k + g(a) + k = g(a) + g(b).$$

The case $a \in L', b \in L^{\downarrow}$.
Analogous to the above.

The case $a, b \in L'$.

Immediate, by the submodularity of f.

Since furthermore $g(z) > k$ whenever $z \in L \setminus L'$, by the definition of M, the function g is as claimed. $\square$

This theorem will also serves as a tool in proving the second theorem, which is the following:

Theorem 8.9. *Let L be a distributive lattice and $L' \subseteq L$ a sublattice. Then there exists a submodular function $f\colon L \to \mathbb{R}_{\geqslant 0}$ and a $k \in \mathbb{R}_{\geqslant 0}$ such that $L' = f^{-1}([0,k))$.*

Theorem 8.8 allows us to first prove Theorem 8.9 only for the special case of sublattices L' which include the top and bottom element of L, and to then handle general sublattices by combing that result with Theorem 8.8.

Lemma 8.10. *Let L be a distributive lattice and $L' \subseteq L$ a sublattice, such that L and L' have the same top and the same bottom element. Then there exists a submodular function $f\colon L \to \mathbb{R}_{\geqslant 0}$ such that $L' = f^{-1}(0)$.*

Proof. By the Birkhoff representation theorem (Theorem 8.1) we may suppose without loss of generality that $L = \mathcal{O}(P)$, for some poset P. We may thus interpret the elements of L (and thus also those of L') as subsets of P.

For every element $p \in P$ let E_p be the set of elements of L' which contain p. In particular, the top element of L lies in E_p, so E_p is non-empty. Thus, we can consider, for every $p \in P$, the set X_p given by $\bigcap_{X \in E_p} X$. Note that p is an element of X_p.

Observe that, since L' is a sublattice, we have $X_p \in L'$ for every p. Given some $Y \in L$ we define $f(Y)$ by summing, over all p in Y, the number of elements of X_p that do not lie in Y. Formally,

$$f(Y) = \sum_{p \in Y} \left| X_p \setminus Y \right| .$$

This function is submodular, since for all $X, Y \in L$ we can calculate as follows

$$\begin{aligned}
&f(X) + f(Y) = \sum_{p \in Y} \left| X_p \setminus Y \right| + \sum_{p \in X} \left| X_p \setminus X \right| \\
&= \sum_{p \in X \cap Y} \left(\left| X_p \setminus Y \right| + \left| X_p \setminus X \right| \right) + \sum_{p \in Y \setminus X} \left| X_p \setminus Y \right| + \sum_{p \in X \setminus Y} \left| X_p \setminus X \right| \\
&= \sum_{p \in X \cap Y} \left(\left| X_p \setminus (X \cap Y) \right| + \left| X_p \setminus (X \cup Y) \right| \right) + \sum_{p \in Y \setminus X} \left| X_p \setminus Y \right| + \sum_{p \in X \setminus Y} \left| X_p \setminus X \right| \\
&= f(X \cap Y) + \sum_{p \in X \cap Y} \left| X_p \setminus (X \cup Y) \right| + \sum_{p \in Y \setminus X} \left| X_p \setminus Y \right| + \sum_{p \in X \setminus Y} \left| X_p \setminus X \right| \\
&\geqslant f(X \cap Y) + \sum_{p \in X \cap Y} \left| X_p \setminus (X \cup Y) \right| + \sum_{p \in Y \setminus X} \left| X_p \setminus (X \cup Y) \right| + \sum_{p \in X \setminus Y} \left| X_p \setminus (X \cup Y) \right|
\end{aligned}$$

$$= f(X \cap Y) + \sum_{p \in X \cup Y} \left| X_p \setminus (X \cup Y) \right|$$
$$= f(X \cap Y) + f(X \cup Y).$$

Thus, all that is left to show is that $f(Y) > 0$ for every $Y \in L \setminus L'$. To see this, we observe that, since the bottom element lies in L', any such Y needs to contain some element p. If $X_p \subseteq Y$ for every $p \in Y$, then this would imply that $Y = \bigcup X_p$, contradicting the assumption that $Y \notin L'$. Thus there is some $p \in Y$ such that $X_p \nsubseteq Y$. In particular there needs to be some $q \in X_p$ such that $q \notin Y$, which witnesses that $f(Y) > 0$. □

Combining Lemma 8.10 and Theorem 8.8 results in a proof of Theorem 8.9:

Proof of Theorem 8.9. Let $\perp$ be the bottom element of L' and let $\top$ be the top element of L'. By Lemma 8.10 there is a submodular function f on $L' = [\perp, \top] \subseteq L$ such that $f^{-1}(0) = L'$ and since L' is finite there is some $k \in \mathbb{R}_{\geqslant 0}$ such that $f^{-1}([0,k)) = L'$. Using this f as input in Theorem 8.8 results in the desired submodular function on L. □

From Theorem 8.8 and Theorem 8.9 we now immediately obtain the same results for subuniverses, in the way discussed above:

Theorem 31. *In other words every subuniverse $\vec{U}'$ of a distributive universe $\vec{U}$ is order-induced submodular in $\vec{U}$.*

Proof. We apply Theorem 8.9 to $\vec{U}'$ as a sublattice of $\vec{U}$ to obtain a submodular function f' and $k' \in \mathbb{R}_{\geqslant 0}$ with $\vec{U}' = f'^{-1}([0,k'))$. We now define a symmetric order function f on $\vec{U}$ with $f(\vec{s}) := f'(\vec{s}) + f'(\overleftarrow{s})$. With $k := 2k'$ we have $\vec{U}' = f^{-1}([0,k)) = \vec{S}_{f,k}$, as desired. □

Theorem 32. *Let $\vec{U}$ be a universe of separations and $\vec{U}' = [\vec{x}, \overleftarrow{x}] \subseteq \vec{U}$ a symmetric interval in $\vec{U}$. Suppose that $f\colon \vec{U}' \to \mathbb{R}_{\geqslant 0}$ is a submodular order function on $\vec{U}'$ with maximum value k. Then there exists a submodular order function $g\colon \vec{U} \to \mathbb{R}_{\geqslant 0}$ such that $g(z) = f(z)$ for all $z \in \vec{U}'$ and $g(z) > k$ for all $z \notin \vec{U}'$.*

Proof. We apply Theorem 8.8 to $\vec{U}'$ as an interval in the lattice $\vec{U}$, to obtain a submodular function g' on $\vec{U}$ which agrees with f on $\vec{U}'$. This function need not be symmetric, but we can define $g(\vec{z}) := \frac{g'(\vec{z}) + g'(\overleftarrow{z})}{2}$. Since f is symmetric and g' agrees with f on $\vec{U}'$, also g agrees with f on $\vec{U}$. Moreover g is symmetric. Since g' takes values larger than k outside of $\vec{U}'$, so does g. □

8.5. Submodular decompositions in distributive universes

In this concluding section we consider decompositions of separation systems which are submodular in some universe, asking how such a separation system can be written as the union of proper subsystems which are still submodular. On one hand, we show that each separation systems $\vec{S}$ which is submodular in some distributive universe $\vec{U}$ of separations can be decomposed (although not necessarily disjoint) into at most three strictly smaller, again submodular in $\vec{U}$, separation systems. On the other hand, we will be able to deduce that we can decompose every such separation system into disjoint submodular subsystems, each of which can be embedded into a universe of bipartitions, in which they are again submodular.

The former statement also allows us to lower bound the size of a largest proper submodular subsystem: by the pigeon-hole principle, at least one of these subsystems will have a size of at least $\frac{|S|}{3}$.[4]

However, while this is a decomposition into fewer parts than the ones we will obtain from our theorems, our decompositions will have the advantage that their constituent subsystems are not merely submodular in $\vec{U}$ but 'spanned' in $\vec{S}$: Given a universe $\vec{U}$ of separations and a subsystem $\vec{S} \subseteq \vec{U}$, we say that $\vec{S}' \subseteq \vec{S}$ is a *corner-closed subsytem of $\vec{S}$ (in $\vec{U}$)* if, for all $\vec{s}, \vec{r} \in \vec{S}'$ we have $\vec{s} \vee \vec{r} \in \vec{S}'$ whenever $\vec{s} \vee \vec{r} \in \vec{S}$. In particular, if $\vec{S}$ is submodular in $\vec{U}$, then any corner-closed subsystem $\vec{S}' \subseteq \vec{S}$ is submodular in $\vec{U}$ as well.

We begin by considering the special case of systems of bipartitions. This will later become a subcase in the proof of our general decomposition theorem. The idea applied in the general case will also be similar to the one in the bipartition case. To be able to transfer these techniques we will apply the Birkhoff representation theorem to a universes of separations and investigate how the involution of the universes interacts with this representation. We will state this in the form of an extended Birkhoff theorem for universes of separations.

8.5.1. Decomposition in bipartition universes

Given the universe $\vec{U}$ of bipartitions of some set V and a separation system $\vec{S} \subseteq \vec{U}$ which is submodular in $\vec{U}$, we consider, for some $v, w \in V$, the set

$$\{\,(A,B) \in \vec{S} : \{v,w\} \subseteq A \text{ or } \{v,w\} \subseteq B\,\}.$$

This set forms a corner-closed subsystem of $\vec{S}$ in $\vec{U}$. We can utilize this observation to find a decomposition of $\vec{S}$ into three proper subsystems.

[4]This observation also links the question of submodular decompositions to the unravelling problem, which we will discuss in the next chapter: suppose $\vec{S}$ contains a separation $\vec{s}$ such that $\vec{S}' = \vec{S} \setminus \{\,\vec{s}, \overleftarrow{s}\,\}$ is still submodular – this is the case if $\vec{S}$ can be unravelled – then we can decompose $\vec{S}$ into the two submodular subsystems $\vec{S}'$ and $\{\,\vec{s}, \overleftarrow{s}, \vec{s} \vee \overleftarrow{s}, \vec{s} \wedge \overleftarrow{s}\,\}$.

Theorem 8.11. *Given a universe $\vec{U} = \vec{\mathcal{B}}(V)$ of bipartitions and a separation system $\vec{S} \subseteq \vec{U}$, such that $|S| \geqslant 3$, there are corner-closed subsystems $\vec{S}_1, \vec{S}_2, \vec{S}_3 \subsetneq \vec{S}$, such that $\vec{S}_1 \cup \vec{S}_2 \cup \vec{S}_3 = \vec{S}$.*

Proof. As $|S| \geqslant 3$, there are two distinct separations $\{A, B\}$, $\{C, D\} \in S$ such that $A, B, C, D \neq \emptyset$. Moreover, we may assume that, after possibly exchanging C and D, we have neither $C \subseteq A$ nor $C \subseteq B$ and thus $A \cap C \neq \emptyset$ and $B \cap C \neq \emptyset$. Additionally, after possibly exchanging A and B, we may assume $B \cap D \neq \emptyset$.

Now pick $x \in A \cap C$, $y \in B \cap C$ and $z \in B \cap D$. Let $\vec{S}_1$ be the set of all separations in $\vec{S}$ not separating x from y, let $\vec{S}_2$ be the set of all separations in $\vec{S}$ not separating x from z and let $\vec{S}_3$ consists of all separations not separating y from z. By construction, theses sets form corner-closed subsystems: a corner of two separations not separating x from y, say, does not separate these two points either.

Moreover, (A, B) is in neither $\vec{S}_1$ nor $\vec{S}_2$ and (C, D) neither in $\vec{S}_2$ nor $\vec{S}_3$, thus $\vec{S}_i \subsetneq \vec{S}$ for all $1 \leqslant i \leqslant 3$.

Finally, observe that, given any $(E, F) \in \vec{S}$, either E or F contains two of the points x, y and z, so $(E, F) \in \vec{S}_1 \cup \vec{S}_2 \cup \vec{S}_3$. Thus $\vec{S}_1 \cup \vec{S}_2 \cup \vec{S}_3 = \vec{S}$, as claimed. □

8.5.2. Birkhoff's theorem for distributive universes and decompositions in distributive universes

To lift Theorem 8.11 to general distributive universes of separations, we will represent separations as subsets of some ground set. For this we will once more, as in Section 8.4, use the Birkhoff representation theorem for distributive lattices:

⇝ p. 170 **Theorem 8.1** (Birkhoff representation theorem; cf. [16, 5.12])**.** *Let L be a finite distributive lattice. The map $\eta\colon L \to \mathcal{O}(\mathcal{J}(L))$ defined by $\eta(a) = \{\, x \in \mathcal{J}(L) : x \leqslant a \,\} = \lceil a \rceil_{\mathcal{J}(L)}$ is an isomorphism of lattices.*

If, in this theorem, the provided distributive lattice L is actually a universe of separations, we obtain an order-reversing involution on $\mathcal{O}(\mathcal{J}(L))$ by concatenating η with the involution on the universe. For our version of the Birkhoff theorem in distributive universes, we examine how this involution behaves with respect to $\mathcal{J}(L)$.

Theorem 33 (Birkhoff representation of universes of separations)**.**
For every involution poset[5] $(P, \leqslant, ')$, the lattice $\mathcal{O}(P)$ becomes a distributive universe of separations $(\mathcal{O}(P), {}^*)$ when equipped with the involution ${}^*\colon X \mapsto P \setminus X'$, where $X' = \{\, x' : x \in X \,\}$.

Let U be a finite distributive universe of separations and let $P = \mathcal{J}(\vec{U})$. Then there exists an order-reversing involution $'$ on P, such that the map $\eta\colon \vec{U} \to \mathcal{O}(P)$ defined by $\eta(a) = \{\, x \in P : x \leqslant a \,\} = \lceil a \rceil_P$ is an isomorphism of universes of separations between $\vec{U}$ and $(\mathcal{O}(P), {}^*)$.

Proof of Theorem 33. The first statement is immediate. For the second part let us assume we are given a distributive universe $\vec{U}$ of separations and need to construct an involution on $P := \mathcal{J}(\vec{U})$ so that $\vec{U}$ is isomorphic to $\mathcal{O}(P)$.

Theorem 8.1 tells us that the two are isomorphic as lattices, so it remains to take care of the involution. Concatenating the isomorphism of lattices $\eta \colon \vec{U} \to \mathcal{O}(J(\vec{U}))$ with the involution on $\vec{U}$ gives us an involution * on $\mathcal{O}(P)$ which is order-reversing. Take note that * maps down-closet subsets of P to down-closed subsets of P; it is *not* defined on the elements of P.

That * is order-reversing means that $X \subsetneq Y$ if, and only if, $X^* \supsetneq Y^*$ for all down-closed subsets X, Y of P. Our aim is to define an order-reversing involution $'$ on P so that for all $X \in \mathcal{O}(P)$ we have $X^* = P \setminus \{\, x' : x \in X \,\}$. We begin with the following claim, which is also a necessary condition for this aim to be achievable:

$$\textit{For all } X \in \mathcal{O}(P) \textit{ we have that } |X^*| = |P| - |X|. \tag{※}$$

We prove (※) by contradiction: assume that X is an inclusion-wise minimal down-closed subset of P for which (※) does not hold. (It clearly holds for the empty set.) Take a maximal element x of X and consider the down-closed set $X - x$. By choice of X, we have $|(X-x)^*| = |P| - |(X-x)|$. From $X^* \subsetneq (X-x)^*$ it thus follows that $|X^*| \leqslant |P| - |X|$.

To see that this holds with equality, first observe that since there is no down-closed set Y with $(X - x) \subsetneq Y \subsetneq X$ and neither is there a down-closed set Y^* with $(X-x)^* \supsetneq Y^* \supsetneq X^*$. However, if $(X-x)^* \setminus X^*$ had more than one element, then adding a minimal one among them to X^* would give such a set Y^*. Hence X^* must be exactly one element smaller than $(X-x)^*$, giving equality and contradicting the choice of X. This proves (※).

Let us now define the involution $'$ on P. The following up- and down-closures are all to be taken in P. For each $x \in P$ we define x' to be the unique element of $(\lceil x \rceil - x)^* \setminus \lceil x \rceil^*$; this is well-defined by (※). We will need to show that $'$ is an involution, that $'$ is order-reversing and that $X^* = P \setminus \{\, x' : x \in X \,\}$ for every down-closed set X.

We have $\lceil x \rceil^* \subseteq P \setminus \lfloor x' \rfloor$, and hence $(P \setminus \lfloor x' \rfloor)^* \subseteq \lceil x \rceil$. If we had proper inclusion, i.e., $(P \setminus \lfloor x' \rfloor)^* \subsetneq \lceil x \rceil$, then the down-closedness of $(P \setminus \lfloor x' \rfloor)^*$ would imply that $(P \setminus \lfloor x' \rfloor)^* \subseteq \lceil x \rceil - x$ and thus $(\lceil x \rceil - x)^* \subseteq P \setminus \lfloor x' \rfloor$, contradicting the choice of x'. Thus the inclusion holds with equality, and we have $\lceil x \rceil^* = P \setminus \lfloor x' \rfloor$.

We are now going to show, given some down-closed set X in which x is maximal, that $(X - x)^* \setminus X^* = \{\, x' \,\}$. Since $\lceil x \rceil \subseteq X$, we have that $X^* \subseteq \lceil x \rceil^*$ and thus X^* cannot contain x'. But $(X-x)^*$ does contain x', as otherwise, by $\lceil x \rceil^* = P \setminus \lfloor x' \rfloor$, we have that $(X-x)^* \subseteq \lceil x \rceil^*$ and thus $(X - x) \supseteq \lceil x \rceil$, which is absurd.

[5] Recall that involution posets are the same as separation systems. However, to emphasise that the involution on $\mathcal{J}(\vec{U})$ is different from the involution on $\vec{U}$, despite $\mathcal{J}(\vec{U})$ being a subset of $\vec{U}$, we prefer the term 'involution poset' in this context.

This observation allows us to infer that $'$ is indeed an involution on P: by the fact that $\lceil x\rceil^* = (\lceil x\rceil - x)^* - x'$ is down-closed, we know that x' is maximal in $(\lceil x\rceil - x)^*$ and x'' is the unique element of $((\lceil x\rceil - x)^* - x')^* \setminus (\lceil x\rceil - x)^{**} = \lceil x\rceil \setminus (\lceil x\rceil - x)$, so x'' is x.

Let us show that we have $X^* = P \setminus \{\, x' : x \in X \,\}$ for all $X \in \mathcal{O}(P)$. We do so by induction on the size of X; for the empty set the statement is immediate. So suppose that the assertion holds for each proper down-closed subset of some non-empty $X \in \mathcal{O}(P)$ and let x be a maximal element of X. Then $(X-x)^* = P \setminus \{\, y' : y \in (X-x) \,\}$. By the earlier observation, the single element in $(X-x)^* \setminus X^*$ is precisely x', giving $X^* = P \setminus \{\, y' : y \in X \,\}$ as claimed.

Finally, we shall check that $'$ is order-reversing. For this let some $x \in P$ be given. Since $\lceil x\rceil^*$ is a down-closed set which does not contain x' we have $\lceil x\rceil^* \subseteq P \setminus \lfloor x'\rfloor$. By applying * to both sides and using the above paragraph we get that $\lceil x\rceil \supseteq P \setminus \{\, y' : y \in P \setminus \lfloor x'\rfloor \,\}$. The right-hand side simplifies to $\{\, y' : y \in \lfloor x'\rfloor \,\}$. Since this set is down-closed and contains $x'' = x$, the inclusion is in fact an equality, i.e., $\lceil x\rceil = \{\, y : y' \in \lfloor x'\rfloor \,\}$. From this it follows that $y \leqslant x$ if and only if $y' \geqslant x'$. $\square$

We are now ready to prove the central decomposition theorem, that every sufficiently large separation system which is submodular inside a distributive host universe of separations, can be either decomposed into three disjoint submodular subsystems, or is isomorphic to a subsystem of a universe of bipartitions while preserving existing corners (i.e., joins and meets). Such an isomorphism $\iota\colon \vec{S} \to \vec{S}'$ between two subsystems $\vec{S} \subseteq \vec{U}$ and $\vec{S}' \subseteq \vec{U}'$ of universes $\vec{U}$ and $\vec{U}'$, where $\iota(\vec{r}) \vee \iota(\vec{s}) = \iota(\vec{r} \vee \vec{s})$ whenever $\vec{r} \vee \vec{s} \in \vec{S}$, and conversely $\iota(\vec{r}) \wedge \iota(\vec{s}) = \iota(\vec{r} \wedge \vec{s})$ whenever $\vec{r} \wedge \vec{s} \in \vec{S}$, for all $\vec{r}, \vec{s} \in \vec{S}$, is called a *corner-faithful embedding.*

Theorem 34. *Let $\vec{U}$ be a distributive universe of separations and let $\vec{S} \subseteq \vec{U}$, $|S| \geqslant 3$, be a separation system which is submodular in $\vec{U}$. Then there are corner-closed subsystems $\vec{S}_1, \vec{S}_2, \vec{S}_3 \subsetneq \vec{S}$ which are submodular in $\vec{U}$ and such that $\vec{S}_1 \cup \vec{S}_2 \cup \vec{S}_3 = \vec{S}$.*

Moreover $\vec{S}_1, \vec{S}_2, \vec{S}_3$ can be chosen disjointly unless $\vec{S}$ can be corner-faithfully embedded into a universe of bipartitions.

Proof. The proof goes via induction on $|U|$. By applying Theorem 33 we may assume, without loss of generality, that $\vec{U} = \mathcal{O}(P)$ for some involution poset $(P, \leqslant, ')$. For every $p \in P$ consider the sets

$$\begin{aligned} \vec{S}_p &:= \{\, X \in \vec{S} : p \in X,\ p' \notin X \,\}, \\ \vec{S}_{p'} &:= \{\, X \in \vec{S} : p \notin X,\ p' \in X \,\}, \\ \vec{S}_{p,p'} &:= \vec{S} \setminus (\vec{S}_p \cup \vec{S}_{p'}). \end{aligned}$$

Note that these are pairwise disjoint, closed under involution, corner-closed and $\vec{S} = \vec{S}_{p,p'} \cup \vec{S}_p \cup \vec{S}_{p'}$. If for any p these three sets form a non-trivial decomposition, we are done. Otherwise either for every $p \in P$ we have $\vec{S} = \vec{S}_{p,p'}$ or for some p we have $\vec{S} = \vec{S}_p$.

If for some p we have $\vec{S} = \vec{S}_p$, then we can consider $\vec{S}$ as a subsystem of $\vec{U}' := \mathcal{O}(P \setminus \{ p, p' \})$ under the corner-faithful embedding $\iota \colon \vec{S}_p \to \vec{U}'$, $X \mapsto X - p$. Since $|U'| < |U|$ we can then apply the induction hypothesis to get the desired decomposition.

If $\vec{S} = \vec{S}_{p,p'}$ for every $p \in P$, then this means, that for every p we have $p \in X \Leftrightarrow p' \in X$ for all $X \in \vec{S}$. In particular, for every X, we have $X^* = X \setminus A' = X \setminus A$. This means that $\vec{S}$ is a submodular subsystem of the bipartition universe $\mathcal{B}(P)$, and Theorem 8.11 gives the desired decomposition. □

Observe that in the universe $\mathcal{O}(P)$ we have $X \wedge X^* = \{ p \in P : p \in X,\, p' \notin X \}$. Hence, when recursively applying the decomposition into $\vec{S}_p$, $\vec{S}_{p'}$ and $\vec{S}_{p,p'}$ as above, we never separate any X and Y where $X \wedge X^* = Y \wedge Y^*$.

Conversely, given any $X \in \mathcal{O}(P)$, the set of all $Y \in \vec{S}$ with $Y \wedge Y^* = X \wedge X^*$ is a corner-closed subsystem of $\vec{S}$. By the last argument of the proof above, these can be considered as subsystems of bipartition universes. We thus obtain our second decomposition result, while also explicitly specifying the subsystems that make up our decomposition:

Theorem 35. *Every separation system $\vec{S}$ which is submodular in some distributive universe $\vec{U}$ of separations is a disjoint union of corner-closed subsystems $\vec{S}_1, \dots, \vec{S}_n$ of $\vec{S}$ (which are thus also submodular in $\vec{U}$) each of which can be corner-faithfully embedded into a universe of bipartitions.*

Specifically, these subsystems are the equivalence classes of the relation $\sim$ on $\vec{S}$ where $\vec{s} \sim \vec{t}$ if and only if $\vec{s} \wedge \overleftarrow{s} = \vec{t} \wedge \overleftarrow{t}$ in $\vec{U}$.

8.6. A separation nested with many others

How nested are the separations in a submodular separation system? This is only a vague question, but we know that reasonably many separations must be nested, since submodularity make the tree-of-tangles theorem work, so a submodular separation system contains a set of at least as many separations as it has profiles, minus one.

Another approach to finding many nested separations is this: using Ramsey's theorem one can easily deduce that every sufficiently large separation system contains either a large set of pairwise nested separations or a large set of pairwise crossing separations. If the separation system is submodular, then for such a large set of pairwise crossing separations we have many pairwise corner separations, which might again be used to construct a large nested set, were it not for the fact that many of the corners might coincide.

As a (non-distributive) example, consider a universe of separations which consists of some large number n of pairwise crossing separations plus just one separation whose orientations form the bottom and top element of our lattice. This universe contains a set of n pairwise crossing separations, yet the largest nested sets (of unoriented separations) consist of only two elements. It does however contain a single separation which is nested with many others.

In the case of a full universe of separations this is hardly surprising, after all the top/bottom separation is always nested with all separations in the universe. Yet, we can show also that every submodular separation system contains a separation which is nested with many others:

Theorem 8.12. *Let $\vec{S}$ be a submodular separation system. Then there exists a separation $s \in S$ which is nested with at least $\sqrt{|S|}$ many separations in S (including s itself).*

Proof. Let $s \in S$ be nested with as many separations from S as possible. Let $N \subseteq S$ denote the set of separations which are nested with s. There are $|S \setminus N| = |S| - |N|$ many separations in S which cross s. Since S is submodular it contains, for every $t \in S$ which crosses s, a corner of s and t; this corner is in N and distinct from s. So, every $t \in S$ which crosses s is nested with at least one element of $N \setminus \{s\}$.

By the pigeon hole principle there then exists an $r \in N$ which is nested with at least $\frac{|S|-b}{b}$ elements of $N \setminus \{s\}$. Since r is also nested with s, it is nested with at least $\frac{|S|-b}{b} + 1$ elements of N. By the choice of b we conclude:

$$b \geqslant \frac{|S| - b}{b} + 1 \implies b^2 \geqslant (|S| - b) + b \implies b \geqslant \sqrt{|S|}\,. \qquad \square$$

9. Submodularity and the unravelling problem

In this chapter we analyse a simple combinatorial problem that is derived from submodularity. The contents of this chapter are based on the preprint [39] which is joint work together with Christian Elbracht and Jakob Kneip.

Here is an intriguingly simple combinatorial problem – simple enough that you can explain it to a first-year student of mathematics – but which is challenging nonetheless:

Problem 9.1 (Unravelling problem)**.** A finite set $\mathcal{X}$ of finite sets is *woven* if, for all $X, Y \in \mathcal{X}$, at least one of $X \cup Y$ and $X \cap Y$ is in $\mathcal{X}$. Let $\mathcal{X}$ be a non-empty woven set. Does there exist an $X \in \mathcal{X}$ for which $\mathcal{X} - X$ is again woven?

Given a set of subsets $\mathcal{X}$ which is woven, an *unravelling* of $\mathcal{X}$ shall be a sequence $\mathcal{X} = \mathcal{X}_n \supseteq \cdots \supseteq \mathcal{X}_0 = \emptyset$ of sets, all woven, such that $|\mathcal{X}_i \setminus \mathcal{X}_{i-1}| = 1$ for all $1 \leqslant i \leqslant n$. If the unravelling problem has a general positive answer, then every woven set will have an unravelling.

The question whether every woven set has an unravelling arose naturally in our study of structurally submodular separation systems. These systems are a framework developed by Diestel, Erde, and Weißauer [25] to generalize the theory of tangles in graphs to an abstract setting, allowing the application of tangles in a multitude of contexts including graphs and matroids, but also other combinatorial structures.

In this chapter we analyse the unravelling problem. We prove affirmative results for versions in two important cases, which come from the original context where the problem arose: submodular separation systems. Our first main result is that unravellings exist for sets $\mathcal{X}$ that consist, for some submodular function f on the subsets of $V = \bigcup \mathcal{X}$, precisely of the sets $X \subseteq V$ with $f(X) < k$ for some integer k. Our second main result settles the unravelling problem for general finite posets, which we call *woven* if they contain, for every two elements, either an infimum or a supremum of these two elements.

In Section 9.1 we give an equivalent formulation of the unravelling problem in terms of distributive lattices and also establish a kind of converse of unravelling, showing that we can find, for every woven set $\mathcal{X}$, some subset of $\bigcup \mathcal{X}$ which we can add to $\mathcal{X}$ and remain woven. We then explain how the unravelling problem turns up in the context of submodularity in abstract separation systems. Throughout the chapter, we will come back to the context of abstract separation systems, to discuss how our results apply there.

In Section 9.2 we give a partial solution to the unravelling problem, by showing that the following class of woven sets, which is particularly important in the theory of tangles, can indeed be unravelled: Let $\mathcal{X}$ be a collection of subsets of some finite set V. If $\mathcal{X}$ has the form $\mathcal{X} = \{ X \in 2^V : f(X) < k \}$ for some function $f\colon 2^V \to \mathbb{R}$ and $k \in \mathbb{R}$, let us say that f *induces* $\mathcal{X}$.

⇝ p. 194 **Theorem 37.** *If $\mathcal{X} \subseteq 2^V$ is induced by a submodular function on 2^V, then $\mathcal{X}$ can be unravelled.*

In Section 9.3 we introduce a possible generalisation of the unravelling problem from subsets of some power set to subsets of any lattice. We show that the lattice analogue of the unravelling problem can be answered in the negative for non-distributive lattices, constructing an explicit counterexample. However, if we restrict this generalized formulation of the unravelling problem to distributive lattices, it becomes equivalent to Problem 9.1.

We conclude in Section 9.4 with our second main result, a variant of the unravelling problem for general partially ordered sets. Let us call a finite poset P *woven* if there exists, for any $p, q \in P$, either a supremum or an infimium in P. A sequence $P = P_n \supseteq \cdots \supseteq P_0 = \emptyset$ of subposets is an *unravelling* of P if P_i is woven and $|P_i \setminus P_{i-1}| = 1$ for every $1 \leqslant i \leqslant n$. Our second main result is that all woven posets have an unravelling:

⇝ p. 202 **Theorem 40.** *Every woven poset can be unravelled.*

Wovenness in posets corresponds to the most general notion of submodularity for separation systems, which we discussed in Section 8.2.

9.1. Unravelling in lattices and separation systems

We can formulate a problem equivalent to Problem 9.1 in terms of lattices. This problem might be easier to work with, and will also allow us to explain how the unravelling problem originally came up:

Problem 9.2. Let L be a finite lattice and $P \subseteq L$ a non-empty woven subset. Does there exist $p \in P$ for which $P - p$ is again woven?

For distributive lattices, Problem 9.2 is equivalent to Problem 9.1 by Birkhoff's representation theorem (see Theorem 8.1), which says that every finite distributive lattice is isomorphic to a sublattice of the subset lattice of some finite set. For general lattices, however, we have a negative solution to Problem 9.2: in Section 9.3 we shall construct a (non-distributive) counterexample for Problem 9.2.

Perhaps surprisingly, it is easy to establish a kind of converse to Problem 9.2: given a lattice L and a woven poset $P \subseteq L$ we can always find a $p \in L \setminus P$ which one can *add* to P while keeping it woven.

Proposition 9.3. *If L is a lattice and $P \subsetneq L$ a proper woven subset of L, then there exists a $p \in L \setminus P$ such that $P + p$ is again woven.*

Proof. Let p be a maximal element of $L \setminus P$. Then $P' := P + r$ is woven: for each $q \in P'$ we have $(p \vee q) \in P'$ by the maximality of p in $L \setminus P$. $\square$

In terms of woven sets in the sense of Problem 9.1, this statement directly implies the following:

Corollary 9.4. *If V is a finite set and $\mathcal{X} \subsetneq 2^V$ woven, then there is a $X \subseteq V$ such that $X \notin \mathcal{X}$ and such that $\mathcal{X} + X$ is again woven.*

The original motivation for considering the unravelling problem originates in submodular separation systems: Given a submodular separation system $\vec{S}$ inside a universe $\vec{U}$ of separations, it might be possible that we find a separation $\vec{s}$ inside $\vec{S}$ which we can delete, together with its inverse, and be left with a separation system $\vec{S} \setminus \{\vec{s}, \overleftarrow{s}\}$ that is again submodular in $\vec{U}$. Formally, given a submodular separation system $\vec{S}$ inside a universe $\vec{U}$ of separations, we are interested if the following property holds:

Property 9.5. There is an $\vec{s} \in S$ such that $\vec{S} \setminus \{\vec{s}, \overleftarrow{s}\}$ is submodular in $\vec{U}$.

If this were to hold for all separation systems which are submodular in $\vec{U}$, then we could recursively apply this reduction step to *unravel* such a separation system, i.e., we would obtain a sequence $\emptyset = \vec{S}_1 \subseteq \vec{S}_2 \cdots \subseteq \vec{S}_n = \vec{S}$ of submodular subsystems such that, for every $i < n$, we have that $\vec{S}_{i+1} \setminus \vec{S}_i$ consists of just one separation $\vec{s}_i$ together with its inverse. Such an unravelling sequence would be of particular use for proving theorems about submodular separation systems via induction. For example, it is possible to obtain a short proof of a tree-of-tangles theorem for submodular separation systems, effectively Theorem 5.4, which have such an unravelling sequence [58, Section 4.1.8].

This question, whether Property 9.5 holds for every structurally submodular separation system, is now closely related to Problem 9.2. In fact, if we could unravel every structurally submodular separation system, we could answer Problem 9.2 positively: If there exists a woven poset P inside a lattice L, such that $P - p$ is not woven, we could construct a structurally submodular separation system inside a universe $\vec{U}$ of separations which can not be unravelled. We use such a construction in Section 9.3 to turn our counterexample to Problem 9.2 into an example of a structurally submodular separation system inside a non-distributive lattice which cannot be unravelled.

Also, the converse of Problem 9.2 established by Proposition 9.3 directly translates to a similar statement about structurally submodular separation systems inside a universe of separations.

Corollary 9.6. *If U is a universe of separations and $S \subsetneq U$ submodular in U, then so is $S + r$ for some $r \in U \setminus S$.*

Proof. Let $\vec{r}$ be a maximal element of $\vec{U} \setminus \vec{S}$. By Proposition 9.3, the separation system $S' := S + r$ is again submodular in U. □

9.2. Unravelling order-induced submodular sets

In this section we show that for a subclass of the woven subsets of a lattice we indeed have unravellings. For this, recall from Chapter 8 that a set P inside a lattice L is *order-induced submodular* if there exists a submodular function $f\colon L \to \mathbb{R}_{\geqslant 0}$ and a real number k such that $P = \{\, p \in L : f(p) < k \,\}$. Here, f being submodular should mean that $f(p) + f(q) \geqslant f(p \vee q) + f(p \wedge q)$ for any $p, q \in L$. Note that every order-induced set P is woven, as the submodularity of f implies that at least one of $f(p \vee q), f(p \wedge q)$ is at most $\max\{\, f(p), f(q) \,\}$ and thus at least one of $p \vee q$ and $p \wedge q$ lies in P, whenever both p and q lie in P. However, there do exists woven sets which are not order-induced submodular, see [37].

We will see in what follows that for order-induced submodular subposets P of a lattice L it is possible to find an *unravelling*, that is a sequence $P = P_n \supsetneq \cdots \supset P_0 = \emptyset$ of posets which are woven in L such that $|P_i \setminus P_{i-1}| = 1$ for every $1 \leqslant i \leqslant n$.

We say that P can be *unravelled* if there exists an unravelling for P. In other words P can be unravelled if we are able to successively delete elements from P until we reach the empty set and maintain the property of being woven throughout.

We shall demonstrate that every order-induced submodular subposet of a lattice can be unravelled.

Theorem 36. *Let L be a lattice with a submodular function f and consider the subset $P = \{\, p \in L : f(p) < k \,\}$ for some k. Then P can be unravelled.*

For the remainder of this section let L be a lattice with a submodular order function f and $P \subseteq L$. It is easy to see that we can perform the first step of an unravelling sequence:

Lemma 9.7. *If $P = \{\, p \in L : f(p) < k \,\}$ and $p \in P$ maximises $f(p)$ in S, then $P - p$ is woven in L.*

Proof. Given $q, r \in P - p$, since P is woven in L at least one of $q \vee r$ and $q \wedge r$ also lies in P. However, by the choice of p we have $f(p) \geqslant f(q)$ and $f(p) \geqslant f(r)$. Thus if one of $q \vee r$ and $q \wedge r$ equals p, the other also needs to lie in P. Thus $P - p$ is indeed woven in L. □

Unfortunately we cannot rely solely on Lemma 9.7 to find an unravelling of P, since after its first application and the deletion of some p the remaining poset $P - p$

may no longer be order-induced submodular. This can happen if $P - p$ contains an r such that $f(r) = f(p)$.

To rectify this, and thereby allow the repeated application of Lemma 9.7, we shall perturb the submodular function f on L to make it injective, whilst maintaining its submodularity and the assertion that $P = \{ p \in L : f(p) < k \}$ for a suitable k. This approach is similar to – and inspired by – the idea of *tie-breaker functions* employed by Robertson and Seymour [64] to construct certain tree-decompositions. For this we show the following:

Theorem 9.8. *Let L be a lattice, then there is an injective submodular function $\rho\colon L \to \mathbb{N}$. Moreover, we can choose ρ so that, for any $p_1, p_2, q_1, q_2 \in L$, we have that $\rho(p_1) + \rho(p_2) = \rho(q_1) + \rho(q_2)$ if and only if $\{ p_1, p_2 \} = \{ q_1, q_2 \}$.*

Proof. Enumerate L as $L = \{ p_1, \dots, p_n \}$. For $q \in L$ let $I(q)$ be the set of all $i \leqslant n$ with $p_i \leqslant q$. We define $\rho\colon L \to \mathbb{N}$ by letting

$$\rho(q) = 3^{n+1} - \sum_{i \in I(q)} 3^i .$$

To see that this function is submodular note that for q and r in L we have $I(q) \cap I(r) = I(q \wedge r)$ and $I(q) \cup I(r) \subseteq I(q \vee r)$. Therefore each $i \leqslant n$ appears in $I(q)$ and $I(r)$ at most as often as it does in $I(q \vee r)$ and $I(q \wedge r)$. This establishes the submodularity.

It remains to show that $\rho(q) \neq \rho(r)$ for all $q \neq r$. For this note that by definition of ρ we have $\rho(q) = \rho(r)$ if and only if $I(q) = I(r)$. But if $q \neq r$, then either $q \notin I(r)$ or $r \notin I(q)$.

To see the moreover part we note that $\rho(p_1) + \rho(p_2) = \rho(q_1) + \rho(q_2)$ if and only if $I(p_1) \cup I(p_2) = I(q_1) \cup I(q_2)$ and $I(p_1) \cap I(p_2) = I(q_1) \cap I(q_2)$. Since $I(p_1)$, $I(p_2)$, $I(q_1)$ and $I(q_2)$ correspond to the down-closures of p_1, p_2, q_1, q_2 in L, this implies that $\{ p_1, p_2 \} = \{ q_1, q_2 \}$: Clearly, if $p_1 = q_1$ then we need to have $p_2 = q_2$, so suppose that $\{ p_1, p_2 \}$ and $\{ q_1, q_2 \}$ are disjoint. Since $p_1 \in I(p_1)$, we see that $p_1 \in I(q_1) \cup I(q_2)$. So suppose without loss of generality that $p_1 < q_1$. Since $q_1 \in I(q_1)$ and $q_1 \notin I(p_1)$, we thus conclude that $q_1 \in I(p_2)$, thus $q_1 < p_2$. As $p_2 \in I(p_2)$, this then implies $p_2 < q_2$. As $q_2 \in I(q_2)$, this is a contradiction as $q_2 \notin I(p_1) \cup I(p_2)$. □

We immediately obtain the following corollary about universes of separations:

Corollary 9.9. *Let U be a universe of separations. Then there is a submodular order function $\gamma\colon \vec{U} \to \mathbb{N}$ with $\gamma(r) \neq \gamma(s)$ for all $r \neq s$.*

Proof. Let ρ be the function obtained from Theorem 9.8 applied to U as a lattice. We set $\gamma(s) = \rho(\vec{s}) + \rho(\overleftarrow{s})$. It is easy to see that this is a submodular order function. The moreover part of Theorem 9.8 guarantees that indeed $\gamma(r) \neq \gamma(s)$ $\forall r \neq s$. □

We can now establish Theorem 36.

Proof of Theorem 36. Let L be a lattice with a submodular order function f. Let $P = \{ p \in L : f(p) < k \}$ for some $k \in \mathbb{R}_{\geqslant 0}$. Let ρ be the submodular function on L from Theorem 9.8. Let ε be the minimal difference between two distinct values of f, that is $|f(p) - f(q)| \geqslant \varepsilon$ or $f(p) = f(q)$ for any two $p, q \in L$. Since L is finite, $\varepsilon > 0$. Pick a positive constant $c \in \mathbb{R}^+$ so that $c \cdot \rho(p) < \varepsilon$ for all $p \in L$. We define a new function $g \colon L \to \mathbb{R}_{\geqslant 0}$ on L by setting

$$g(p) := f(p) + c \cdot \rho(p)\,.$$

Then g is submodular and, like ρ, has the property that $g(p) \neq g(q)$ whenever $p \neq q$. Enumerate the elements $p_1, \dots, p_n$ of P so that $g(p_1) < g(p_2) < \cdots < g(p_n)$. Then $P_i := \{ p_1, \dots, p_i \} \subseteq P$ is woven in L for each $i \leqslant n$: for $i = n$ it equals P, and for $i < n$ we have that $P_i = \{ p \in L : g(p) < g(p_{i+1}) \}$, which is woven in L since g is a submodular function on L. Thus $P = P_n \supseteq \cdots \supseteq P_0 = \emptyset$ is an unravelling for P. □

Theorem 36 allows us to give a class of sets $\mathcal{X} \subseteq 2^V$ for which we can answer Problem 9.1 positively:

Theorem 37. *If $\mathcal{X} \subseteq 2^V$ is induced by a submodular function on 2^V, then $\mathcal{X}$ can be unravelled.*

Proof. By adding a large constant to $f(X)$ for every $X \subseteq V$ we may suppose that $f(X) \geqslant 0\ \forall X \subseteq V$. Applying Theorem 36 to the subset-lattice 2^X together with its subset $\mathcal{X}$ results in the desired sequence $\emptyset = \mathcal{X}_0 \subseteq \mathcal{X}_1 \subseteq \cdots \subseteq \mathcal{X}_n = \mathcal{X}$. □

As a corollary, Theorem 36 also allows us to show that separation systems $\vec{S}_k$ inside a universe of separations with a submodular order function can be unravelled.

Theorem 38. *Let $\vec{U}$ be a universe of separations with a submodular order function f and $\vec{S} = \vec{S}_k$ for some k. Then $\vec{S}$ can be unravelled.*

Proof. Perform the same argument as in the proof of Theorem 36, using the function γ from Corollary 9.9 instead of the function ρ from Theorem 9.8. □

9.3. A woven subset of a lattice which cannot be unravelled

In this section we are going to construct a counterexample to Problem 9.2 for non-distributive lattices. So, we construct a lattice L together with a woven subset P of L so that $P - p$ is not woven in L for any $p \in P$.

This construction needs to be such that for every element p of P there are elements q and r of P such that either $p = q \vee r$ and $q \wedge r \notin P$ or $p = q \wedge r$ and $q \vee r \notin P$.

We will construct our lattice L by building its Hasse diagram. To be able to prove that our construction results in a lattice we need to start with a graph of high girth. Specifically we will use a 4-regular graph of high girth as a starting point. Lazebnik and Ustimenko have constructed such graphs:

Lemma 9.10 ([59]). *There exists a 4-regular graph G with girth at least eleven.*

For contradiction arguments we will try to find short closed walks in our graph. The following simple lemma then tells us that these contradict the high girth of G:

Lemma 9.11. *If G is a graph, $W = v_1 v_2 \dots v_n v_1$ a closed walk in G such that there exists an j with $v_i \neq v_j$ for all $i \neq j$ and $v_{j-1} \neq v_{j+1}$, then W contains a cycle. In particular, G contains a cycle of length at most n.*

Proof. Since $v_j \neq v_i$ for all $i \neq j$, the graph $W - v_j$ is connected. Thus, $W - v_j$ contains a path between v_{j-1} and v_{j+1} which together with v_j forms the desired cycle. □

We are now ready to start the construction of our lattice L together with its woven subset P.

Let G be a 4-regular graph of girth at least eleven. The ground set of our lattice L consists of a top element t, a bottom element b and four disjoint copies of $V(G)$ which we call V^-, V, W and W^+.

We say that $v \in V^- \cup V \cup W \cup W^+$ *corresponds* to $w \in V^- \cup V \cup W \cup W^+$ if they are copies of the same vertex in $V(G)$.

We now start with defining our partial order on L. We define, for $v \in V$ and $w \in W$, that $v \leqslant w$ if and only if there is an edge between v and w in G.

Now consider the bipartite graph G' on $V \cup W$ where $v \in V$ is adjacent to $w \in W$ if and only if $v \leqslant w$. This bipartite graph is 4-regular graph and has girth at least twelve. Every regular bipartite graph has a 1-factor. Hence, G' has a colouring of $E[G']$ with two colours, *red* and *blue* say, such that every vertex in G' is adjacent to exactly two red and exactly two blue edges. We fix one such colouring.

To define our partial order for $v^- \in V^-$ and $v \in V$ we define that $v^- \leqslant v$ if and only if there is a red edge between v and the vertex in W corresponding to v^-. Thus, every v^- in V^- lies below exactly two points in V, we call these the *neighbours in V* of v^-.

Similarly, we let $w \leqslant w^+$ for $w \in W$ and $w^+ \in W^+$ if and only if there is a blue edge between w and the vertex in V corresponding to w^+. We call the two points in W which lie below $w^+ \in W^+$ the *neighbours in W* of w^+.

We finish our definition of $\leqslant$ by taking the transitive closure and defining $b \leqslant v$ and $v \leqslant t$ for every $v \in L$. It is easy to see that this $\leqslant$ is indeed a partial order.

We claim that $(L, \leqslant)$ is a lattice, that $P = V \cup W \cup \{ t, b \} \subseteq L$ is a woven subset of L and that $P - p$ is not woven in L for any $p \in P$. To show that L is a lattice and that P is woven in L we have to show that there is, for every pair $x, y \in L$, a supremum and an infimum and that at least one of these two lies in P if $x, y \in P$. We do so via a series of lemmas which distinguish different cases for x, y.

Let us first consider the case that either both x and y lie in V, or that they both lie in W:

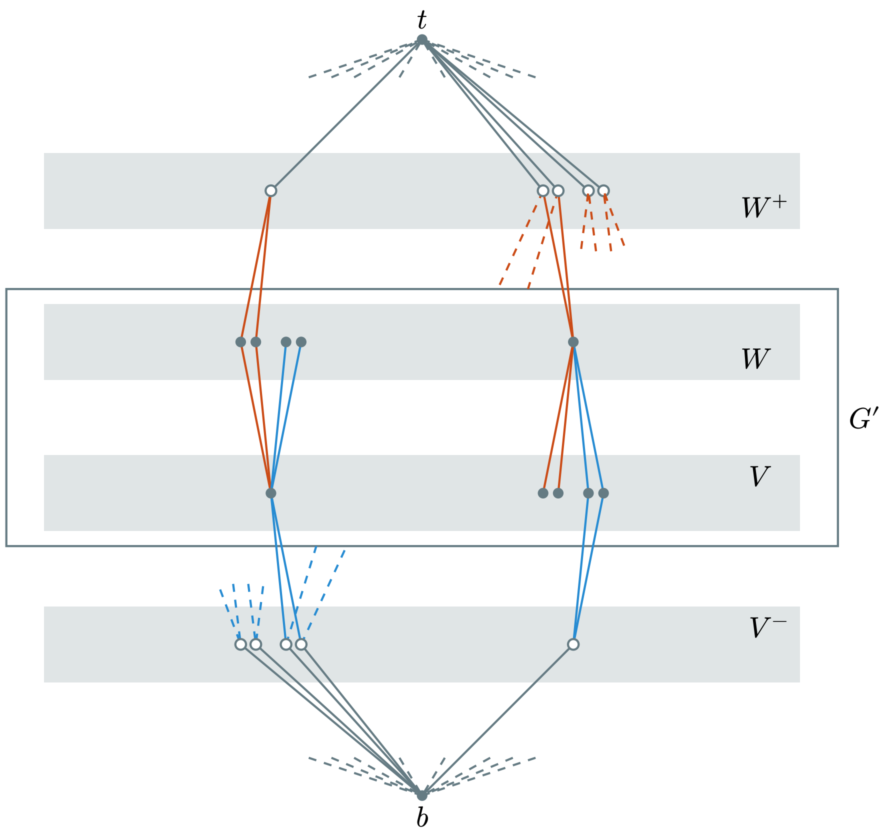

Figure 9.1.: The Hasse diagram of L. The points in P are denote by black dots, the points outside of P are white.

Lemma 9.12. *If $v_1, v_2 \in V$, then there is a supremum and an infimum of v_1, v_2 in L. Moreover, if $v_1 \wedge v_2 \neq b$ then $v_1 \vee v_2 \in W$.*

Analogously, if $w_1, w_2 \in W$, then there is a supremum and an infimum of w_1, w_2 in L. Moreover, if $w_1 \vee w_2 \neq t$ then $w_1 \wedge w_2 \in V$.

Proof. Let us start by showing that there is a supremum of v_1 and v_2.

First consider the case that the neighbourhoods of v_1 and v_2 in G' intersect, that is, $N_{G'}(v_1) \cap N_{G'}(v_2) \neq \emptyset$. In this case, there is only one point in the intersection, since if there are $w_1, w_2 \in N_{G'}(v_1) \cap N_{G'}(v_2), w_1 \neq w_2$, then $v_1 w_1 v_2 w_2 v_1$ would be a cycle of length 4 in G', contradicting the fact that G' has girth at least 12. We claim that the single point in the intersection, which we call w, is the supremum of v_1 and v_2.

To see this consider any $x \in L$ such that $v_1 \leqslant x, v_2 \leqslant x$. We need to show that $w \leqslant x$. If $x = t$ then this is clear and $x \in W \cup V \cup V^- \cup \{b\}$ is not possible, so suppose that $x \in W^+$. Let w_1, w_2 be the neighbours in W of x, i.e., $w_1, w_2 \leqslant x$.

We show that $w_1 = w$ or $w_2 = w$. So suppose that $w \neq w_1, w_2$. Let $v_x \in V$ be the point corresponding to x. As $v_1 \leqslant x$ we may suppose without loss of generality that $v_1 \leqslant w_1$. Now if $v_2 \leqslant w_2$ then $wv_1w_1v_sw_2v_2w$ contains a cycle of length at most 6 in G' by Lemma 9.11, as $v_1 \neq v_2$ and $w \notin \{ v_1, w_1, v_s, w_2, v_2 \}$. This contradicts the fact that G' has girth at least 12. Thus, $v_2 \leqslant w_1$ and, hence, $w_1 = w$ as $N_{G'}(v_1) \cap N_{G'}(v_2) = \{w\}$, contradicting the assumption that $w \neq w_1$ and thus proving $w \leqslant x$.

Now suppose that $N_{G'}(v_1) \cap N_{G'}(v_2) = \emptyset$.

Then every candidate for a supremum of v_1 and v_2 is either t, or lies in W^+, hence it is enough to show that there cannot be two elements $w_1^+, w_2^+ \in W^+$ both satisfying $v_1, v_2 \leqslant w_1^+, w_2^+$. So suppose that there are two such points and denote the neighbours of w_1^+ and w_2^+ in W as w_{11}, w_{12} and w_{21}, w_{22} respectively, i.e., $w_{11}, w_{12} \leqslant w_1^+$ and $w_{21}, w_{22} \leqslant w_2^+$.

As $v_1 \leqslant w_1^+, w_2^+$, we may suppose without loss of generality that $v_1 \leqslant w_{11}, w_{21}$. Since $N_{G'}(v_1) \cap N_{G'}(v_2) = \emptyset$, we thus have $v_2 \leqslant w_{12}, w_{22}$ and $w_{12}, w_{22} \notin \{ w_{11}, w_{21} \}$. Let us denote the corresponding points of w_1^+ and w_2^+ in V as $v_{w_1^+}$ and $v_{w_2^+}$. Since $w_1^+ \neq w_2^+$ either $w_{12} \neq w_{22}$ or $w_{11} \neq w_{21}$, as otherwise G' would contain a cycle of length four. In any case, we consider the closed walk $v_1w_{11}v_{w_1^+}w_{12}v_2w_{22}v_{w_2^+}w_{21}v_1$. As $v_{w_1^+} \neq v_{w_2^+}$, we have $v_1 \neq v_{w_1^+}$ or $v_1 \neq v_{w_2^+}$ and $v_2 \neq v_{w_1^+}$ or $v_2 \neq v_{w_2^+}$. Furthermore, either $w_{11} \notin \{ w_{12}, w_{21}, w_{22} \}$ and $w_{21} \notin \{ w_{11}, w_{12}, w_{22} \}$ or $w_{12} \notin \{ w_{11}, w_{21}, w_{22} \}$ and $w_{22} \notin \{ w_{11}, w_{12}, w_{21} \}$ This allows the application of Lemma 9.11 to our walk, yielding a cycle of length at most 8, which contradicts the fact that G' has girth at least 12. Thus, there exists a supremum $v_1 \vee v_2$ in L.

One candidate for the infimum $v_1 \wedge v_2$ is b. Every other candidate needs to lie in V^-. However, there can be at most one such candidate in V^-, otherwise, these candidates together with v_1, v_2 would correspond to a cycle of length 4 in G' contradicting the fact that G' has girth at least 12. Thus, there is indeed an infimum $v_1 \wedge v_2$.

Moreover, if $v_1 \wedge v_2 \neq b$, then there is a point $w \in W$ such that both, v_1w and v_2w are red edges in G', hence $N_{G'}(v_1) \cap N_{G'}(v_2) \neq \emptyset$, which shows the moreover part of the claim.

The statement for $w_1, w_2 \in W$ follows by a symmetric argument. □

We can now apply Lemma 9.12 to show the existence of suprema and infima between $v \in V$ and $w \in W$:

Lemma 9.13. *If $v \in V$ and $w \in W$, then there is a supremum and an infimum of v and w in L. Moreover, if $v \wedge w \neq b$ then $v \vee w = t$ or $v \leqslant w$.*

Proof. If $v \leqslant w$ then the statement is obvious, so suppose that $v \nleqslant w$.

By Lemma 9.12, every point $w_i \in N_{G'}(v)$ has a supremum with w which is either t or lies in W^+. Moreover, there can be at most one point $w_i \in N_{G'}(v)$ such that the supremum $w_i \vee w$ is in W^+, since if there are two, $w_1, w_2 \in N_{G'}(v)$ say, then, by

Lemma 9.12, $w_1 \wedge w$ and $w_2 \wedge w \in$ both lie in V and thus $wv_1w_1vw_2v_2w$ is a cycle of length 6 in G'. Hence $v \vee w$ is well-defined.

A symmetric argument shows that also $v \wedge w$ is well-defined, so all that is left to show is that $v \vee w \in W^+$ and $v \wedge w \in V^-$ cannot both occur.

However, if this were the case, say $w^+ = v \vee w \in W^+$ and $v^- = v \wedge w \in V^-$, we can consider the corresponding vertex v_{w^+} of w^+ in V and the corresponding vertex w_{v^-} of v^- in W. By definition, there is a vertex $w_1 \in W$ such that $vw_1 \in E(G')$ and both $w_1v_{w^+}$ and wv_{w^+} are blue edges. Similarly, there is a vertex $v_1 \in V$ such that $v_1w \in E(G')$ and both $v_1w_{v^-}$ and vw_{v^-} are red edges. Consider the closed walk $vw_1v_{w^+}wv_1w_{v^-}v$. We have $v \notin \{ v_{w^+}, v_1 \}$ as $v \nleq w$ and similarly $w \notin \{ w_1, w_{v^-} \}$. Moreover, since every edge in G' has precisely one colour we have $v_1w_{v^-} \neq v_{w^+}w_1$ and thus either $w_{v^-} \neq w_1$ or $v_1 \neq v_{w^+}$. We can thus apply Lemma 9.11 to our walk to show the existence of a cycle of length at most 6 in G', which is a contradiction. $\square$

Finally it remains to consider suprema $x \vee y$ and infima $x \wedge y$ where one of x and y lies in V^- or W^+:

Lemma 9.14. *If $v^- \in V^-$ and $x \in L$, then there exists a supremum and an infimum of v and x in L.*

Similarly, if $w^+ \in W^+$ and $x \in L$, then there exists a supremum and an infimum of v and x in L.

Proof. If v^- and x are comparable, the statement is obvious, so suppose that this is not the case. It is then immediate that $v^- \wedge x = b$.

Let v_1, v_2 be the two points in V such that $v^- = v_1 \wedge v_2$ and let w_{v^-} be the point in W corresponding to v^-. We note that any $l \in L$ satisfies $v^- < l$ if and only if $v_1 \leqslant l$ or $v_2 \leqslant l$. We distinguish multiple cases, depending on whether x lies in W^+, W, V or V^-.

If $x \in W^+$, then $x \vee v^- = t$.

If $x \in W$, then $x \vee v_1$ and $x \vee v_2$ exist by Lemma 9.13 and it is enough to show that $x \vee v_1$ and $x \vee v_2$ are comparable. If they are incomparable, then $x \vee v_1 \in W^+$ and $x \vee v_2 \in W^+$ and moreover $x \vee v_1 \neq x \vee v_2$ and $v_1 \nleq x \vee v_2$ as well as $v_2 \nleq x \vee v_1$. Let $v_3 \in V$ be the point corresponding to $x \vee v_1$, let $v_4 \in V$ be the point corresponding to $x \vee v_2$, let $w_3 \in W$ such that $w_3 \vee x = x \vee v_1$ and let $w_4 \in W$ such that $w_4 \vee x = x \vee v_2$. Note that both v_1, v_2, v_3, v_4 and x, w_{v^-}, w_3, w_4 consist of pairwise distinct points as $v_1 \nleq x$ and $v_2 \nleq x$ and $w_{v^-} \notin \{ x, w_3, w_4 \}$, thus $w_3v_3wv_4w_4v_2w_{v^-}v_1w_3$ needs to be a cycle of length 8 in G' contradicting the fact that G' has girth at least 12.

If $x \in V$, then again $x \vee v_1$ and $x \vee v_2$ exist by Lemma 9.13, and if they are incomparable we may suppose that $x \vee v_1 \in W^+ \cup W$ and $x \vee v_2 \in W^+ \cup W$ and moreover $x \vee v_1 \neq x \vee v_2$.

If $x \vee v_1 \in W$ and $x \vee v_2 \in W$, then $v_1w_{v^-}v_2(x \vee v_2)x(x \vee v_1)v_1$ would be a cycle of length 6 in G' as $x \vee v_1 \neq x \vee v_2$.

Now suppose that $x \vee v_1 \in W$ and $x \vee v_2 \in W^+$. Let v_{x_2} be the point in V corresponding to $x \vee v_2$ and let $w_1, w_2 \in W$ such that $w_1 \vee w_2 = x \vee v_2$. We may suppose that $w_1, w_2 \neq x \vee v_1$ and that $v_2 \leqslant w_1$ and $x \leqslant w_2$. Note that $v_{x_2} \neq x$ as otherwise $x \vee v_x = w_2$. Now $xw_2v_{x_2}w_1v_2w_{v^-}v_1(x \vee v_1)x$ contains a cycle of length at most 8 in G' by Lemma 9.11, as $s \notin \{ v_1, v_2, v_{x_2} \}$ and $x \vee v_1 \neq w_2$.

So we may suppose that $x \vee v_1 \in W^+$ and $x \vee v_2 \in W^+$.

Let v_{x_1} be the point in V corresponding to $x \vee v_1$, v_{x_2} be the point in V corresponding to $x \vee v_2$, let $w_1, w_2, w_3, w_4 \in W$ such that $w_1 \vee w_2 = x \vee v_1$ and $w_3 \vee w_4 = x \vee v_2$. We may suppose that $v_1 \leqslant w_1$, $v_2 \leqslant w_3$ and $x \leqslant w_2, w_4$. Note that $x \notin \{ v_1, v_2, v_{x_1}, v_{x_2} \}$ and that $w_4 \neq w_2$ as otherwise $w_4 \leqslant x \vee v_1$ and thus $x \vee v_2 = x \vee w_4 \leqslant x \vee v_1$. Thus, $xw_2v_{x_1}w_1v_1w_{v^-}v_2w_3v_{x_2}w_4x$ contains a cycle in G' of length at most 10 by Lemma 9.11.

So the remaining case is $x \in V^-$. Let us denote the vertex in W corresponding to x as w_x and let $v_3, v_4 \in V$ such that $v_3 \wedge v_4 = x$. Since every candidate for a supremum of v^- and x lies above one of $v_1 \vee v_3, v_1 \vee v_4, v_2 \vee v_3$ and $v_2 \vee v_4$, all of which exist by Lemma 9.13, it is enough to show that all these points are comparable, since then the smallest of them needs to be the supremum of v^- and x.

However, we know by the previous argument that $v^- \vee v_3$ exists, which needs to be equal to $v_1 \vee v_3$ or $v_2 \vee v_3$. Hence $v_1 \vee v_3$ and $v_2 \vee v_3$ are comparable.

Similarly, if we consider $v^- \vee v_4$ we see that $v_1 \vee v_4$ and $v_2 \vee v_4$ are comparable.

If we consider $x \vee v_1$, we observe that $v_1 \vee v_3$ and $v_1 \vee v_4$ are comparable.

And finally, if we consider $x \vee v_2$, we see that $v_2 \vee v_3$ and $v_2 \vee v_4$ are comparable as well and therefore there indeed exists a supremum of v^- and x. □

We have now seen that L is indeed a lattice and that P is woven in L. This allows us to state and prove the main result of this section:

Theorem 9.15. *L is a lattice and $P = V \cup W \cup \{ t, b \} \subseteq L$ is woven in L such that $P - p$ is not woven in L for any $p \in P$.*

Proof. By Lemma 9.12 to 9.14 L is indeed a lattice. To see that P is woven in L observe that by Lemma 9.12, Lemma 9.13 and the fact that t and b are comparable with every element in P it follows that at most one of $x \vee y$ and $x \wedge y$ lie outside of P, for any $x, y \in P$.

For any $p \in V$ there are $w_1, w_2 \in W$ such that pw_1 and pw_2 are both blue edges in G', thus both $w_1 \vee w_2$ and $w_1 \wedge w_2$ lie outside of $P - p$. Similarly, $P - p$ is not woven in L for any $p \in W$. Finally, if $p = b$ we note that there are $v_1, v_2 \in V$ such that $v_1 \vee v_2 \in W^+$ which implies that $v_1 \wedge v_2 = b$ and shows that $P - b$ is not woven in L. Similarly, $P - t$ is not woven in L. □

As before, this result about woven subsets of lattices allows us to directly obtain a result about structurally submodular separation systems, as we can use this lattice L to construct a universe $\vec{U}$ of separations together with a structurally submodular separation system $\vec{S} \subseteq \vec{U}$ which cannot be unravelled:

Theorem 39. *There exists a universe $\vec{U}$ of separations and a submodular subsystem $\vec{S} \subseteq \vec{U}$ such that $\vec{S} \setminus \{\vec{s}, \overleftarrow{s}\}$ is not submodular in $\vec{U}$ for any $\vec{s} \in \vec{S}$.*

Proof. Let L' be a copy of L with reversed partial order, i.e., the poset-dual of L. In the disjoint union $L \sqcup L'$ we now identify the copy of t in L (the top of L) with the copy of b in L' (the top of L') and the copy of b in L with the copy of t in L' to obtain $\vec{U}$. It is easy to see that this forms a universe of separations and that $\vec{S} = P \cup P'$ (where $P \subseteq L$ is as above and $P' \subseteq L'$ is the image of P in L') is a separation system which is submodular in $\vec{U}$. Moreover, there is no separation $\vec{s} \in \vec{S}$ such that $\vec{S} \setminus \{\vec{s}, \overleftarrow{s}\}$ is again submodular in $\vec{U}$. □

Note that neither our lattice L nor the constructed universe $\vec{U}$ of separations are distributive.

9.4. Woven posets

Instead of asking in Problem 9.2 for a woven subset P inside a lattice L, we might as well directly ask for a partially ordered set P, which is woven in itself. More precisely let us say that a partially ordered set P is *woven* if we have, for any two elements p, q of P a supremum or an infimum *in* P, i.e., there exists a $r \in P$ such that $p \leqslant r, q \leqslant r$ and $r \leqslant s$ whenever $s \in P$ such that $q \leqslant s$ and $p \leqslant s$ or there exists a $r \in P$ such that $p \geqslant r, q \geqslant r$ and $r \geqslant s$ whenever $s \in P$ such that $q \geqslant s$ and $p \geqslant s$.

The *Dedekind-MacNeille-completion* [62] from lattice theory implies that we can find, for each poset P, a lattice L in which P can be embedded in such a way that existing joins and meets in P are preserved. Hence, if P is a finite woven set there exists a lattice L in which P can be embedded so that the image of P in L is woven in L.

Using this notion of wovenness inside the poset itself, we can now weaken the concept of unravelling, by considering a woven poset P instead of a woven subset of a lattice. We will be able to show that, given a woven poset P, we can always remove a point so that the remainder is again a woven poset.

Even though every woven poset can be embedded into a lattice, this still is a proper weakening of the unravelling conjecture. The key difference here lies in the different perspective we take on $P - p$, given a poset P and some $p \in P$: if we consider P as a woven poset and $P - p$ is again woven, then there are lattices L and L' in which P and $P - p$, respectively, can be embedded so that the images are woven as subset of these lattice. However, these two lattices are different, and in general it is not possible to find one lattice in which both P and $P - p$ can be embedded so that their images are woven in that lattice. In this sense, having an unravelling for the wovenness of a poset is a weaker property than having an unravelling as a woven subset of a lattice.

To prove this weaker unravelling property for woven posets we will show that every woven poset contains a point p with precisely one lower (or one upper) cover, i.e., there exists precisely one q such that $p > q$ ($p < q$) and there does not exist any $c \in P$ such that $p > c > q$ ($p < c < q$). Deleting such a point does not destroy the wovenness, as shown by the following lemma:

Lemma 9.16. *Let P be a woven poset and $p \in P$ a point with precisely one lower (upper) cover p', then $P' = P - p$ is a woven poset.*

Proof. Let $x, y \in P'$. We need to show that x and y have a supremum or an infimum in P'. If they have a supremum s in P, then $s \neq p$: as p' is the only lower cover of p we have $x, y \leqslant p'$ as soon as $x, y \leqslant p$. Thus $s \in P'$ is also the supremum of x and y in P'.

If x, y have an infimum z in P, then either $z \neq p$ and z is also the infimum in P' or $z = p$, in which case p' is the infimum of x and y in P', as p' is the only lower cover of p.

The upper cover case is dual. $\square$

Thus, what is left to show is that there always exists a point $p \in P$ with precisely one upper or precisely one lower cover. To see this, we consider the maximal elements of P, since any subset of them needs to have an infimum by the following lemma:

Lemma 9.17. *Let P be a woven poset and M its set of maximal elements. Then every non-empty subset $M' \subseteq M$ has an infimum $\mathsf{inf}\, M'$ in P.*

Proof. We proceed by induction on $|M'|$. For the induction start $|M'| = 1$ this is trivial. For the induction step consider $|M'| \geqslant 2$ and let $m \in M'$ and $M'' := M' - m$. By the inductive hypothesis M'' has an infimum p. Since m is maximal there can only be a supremum of m and p if m and p are comparable. However, then there also exists an infimum of m and p in P. Thus, as P is woven, in any case P needs to contain an infimum q of m and p. This q lies below all of M' and, conversely, every point which lies below all of M' lies below both p and m and hence below q. Thus q is the infimum of M' in P. $\square$

Given a woven poset P, let M be the set of maximal elements of P. Given some subset $M' \subseteq M$ we are interested in those points $x \in P$ where, for every maximal element $m \in M$ we have $x \leqslant m$ precisely if $m \in M'$. Let us denote as $d(M')$ the set of all these points in P.

Either each such set $d(M')$ just consist of at most one point, or there is some M' such that $d(M')$ has size more than one. In the latter case, the following lemma guarantees that we find a point $p \in P$ with only one upper cover:

Lemma 9.18. *Let P be a woven poset and M the set of maximal elements of P. If $M' \subseteq M$ is subset-minimal with the property that $d(M')$ contains at least two points, then there is an $x \in d(M')$ for which $\mathsf{inf}\, M'$ is the only upper cover.*

Proof. Observe that, if $d(M) \neq \emptyset$ then $\inf M' \in d(M)$. Let x be a maximal element of $d(M') - \inf M'$. Since x is a candidate for $\inf M'$, we have that $\inf M'$ is an upper cover of x. If y is any point other than $\inf M'$ such that $x < y$ then y lie in $d(M'')$ for some proper subset M'' of M. Thus, by our assumption, y is the only element of $d(M'')$ and therefore $y = \inf M''$. However, $\inf M' \leqslant \inf M''$ and $y \neq \inf M'$, thus y is not an upper cover of x. □

It remains to consider the case where every $d(M')$ has size one. However, in that case we can find an element with only one lower cover, as shown in the following lemma:

Lemma 9.19. *Let P be a woven poset. Then P has an element which has precisely one lower or one upper cover.*

Proof. Suppose the converse is true. Let M be the set of maximal elements of P. Note that every element of P lies in $d(M')$ for exactly one set $M' \subseteq M$. By Lemma 9.18, given any $M' \subseteq M$ there exists at most one element in $d(M')$. Moreover, by Lemma 9.17 we know that $\inf M'$ exists for every $M' \subseteq M$.

Now if $|d(M')| = 1$ for some $M' \subseteq M$, then $\inf M' \in d(M')$: we know that $\inf M'$ is in $d(M'')$ for some $M'' \subseteq M$ and clearly $M' \subseteq M''$, however if $d(M') = \{v\}$, say, then clearly $v \leqslant \inf M'$ which implies that $M'' \subseteq M'$ and thus $M' = M''$.

However, since every element of P lies in some $d(M')$ and $\inf M' \leqslant \inf M''$ whenever $M'' \subseteq M'$ this implies that $\inf M$ is the smallest element of P. However, any upper cover of this smallest element $\inf M$ has precisely one lower cover, which is a contradiction. □

Thus, if we consider woven posets instead of woven subsets of a fixed lattice (as in Section 9.3) we can indeed unravel every such poset: given some woven poset P, by Lemma 9.19, P contains an element p which has only one upper or lower cover, and, by Lemma 9.16, $P - p$ is again woven. We obtain the following theorem:

Theorem 40. *Every woven poset can be unravelled.*

Again we can translate this result to abstract separation systems. Consider separation systems $\vec{S}$ which are submodular on their own, not in the context of a surrounding universe $\vec{U}$ of separations, that is, for any two separations $\vec{s}, \vec{t} \in \vec{S}$ a supremum or an infimum in $\vec{S}$ – this is the same as for woven posets, but only applying Theorem 40 to these does not give us an unravelling of $\vec{S}$, since we always need to delete a separation together with its inverse. We will take care of that, proving the following corollary of Theorem 40:

Theorem 41. *Let $\vec{S}$ be a submodular separation system. Then there exists an $\vec{s} \in \vec{S}$ such that $\vec{S} \setminus \{\vec{s}, \overleftarrow{s}\}$ is again submodular.*

Proof. Observe that $\vec{S}$ considered as a poset is woven. Let M be the set of maximal elements of $\vec{S}$. We note that $\vec{s} \geqslant \overleftarrow{t}$ for all $\vec{s}, \vec{t} \in M$. Therefore $\mathsf{inf}\, M \geqslant \overleftarrow{t}$ for all $\vec{t} \in M$ and thus $\mathsf{inf}\, M \geqslant \mathsf{sup}\, M^* = (\mathsf{inf}\, M)^*$. Suppose that there is a proper subset M' of M such that $|d(M')| \geqslant 2$ and let M' be chosen subset-minimal with that property. Let $\vec{x} \in d(M')$ be as guaranteed by Lemma 9.18. We note that $\vec{x} \neq (\mathsf{inf}\, M')^*$ as otherwise $\vec{x} \leqslant (\mathsf{inf}\, M)^* \leqslant \mathsf{inf}\, M$, contradicting the fact that $\vec{x} \in d(M')$. But this implies that $\vec{S} - \vec{x}$ is a woven poset by Lemma 9.16. However, $\overleftarrow{x}$ has only one lower cover in $\vec{S}$ and, since this cover is not $\vec{x}$, also exactly one lower cover in $\vec{S} - \vec{x}$. Thus, again by Lemma 9.16, also $(\vec{S} - \vec{x}) - \overleftarrow{x}$ is a woven poset and thus $S - x$ is a submodular separation system.

Hence, we may suppose that $|d(M')| \leqslant 1$ for all proper subset M' of M. This implies that every element $\vec{s} \in \vec{S}$ is nested with $\mathsf{inf}\, M$: if $\vec{s} \in d(M)$ then $\vec{s} \leqslant \mathsf{inf}\, M$ and if $\vec{s} \in d(M')$ for a proper subset M' of M, then $\vec{s} = \mathsf{inf}\, M' \geqslant \mathsf{inf}\, M$. Now suppose that $|M| \geqslant 2$. Then there is a $\vec{m} \in M$ such that $\overleftarrow{m} \neq \mathsf{inf}\, M$. We claim that $\vec{S} \setminus \{ \vec{m}, \overleftarrow{m} \}$ is again submodular. To see this suppose that, for some $\vec{x}, \vec{y} \in \vec{S}$, we have that $\vec{x} \vee \vec{y} = \vec{m}$ (the case $\vec{x} \wedge \vec{y} = \overleftarrow{m}$ is dual). As $\vec{x}$ and $\vec{y}$ are nested with $\mathsf{inf}\, M$ this implies that $\vec{x}, \vec{y} \geqslant \mathsf{inf}\, M$ as $\vec{x} \leqslant \mathsf{inf}\, M$ would imply that $\vec{x} \vee \vec{y} = \vec{y}$ or $\vec{x} \vee \vec{y} \leqslant \mathsf{inf}\, M$. Thus, $\vec{x} = \mathsf{inf}\, M'$ and $\vec{y} = \mathsf{inf}\, M''$ for subsets M', M'' of M, say. Thus, $\mathsf{inf}(M' \cup M'')$, which exists by Lemma 9.17, is also the infimum of $\vec{x}$ and $\vec{y}$. Moreover, since $\overleftarrow{m} \neq \mathsf{inf}\, M$ and $\overleftarrow{m}$ is a minimal element of $\vec{S}$ and $\mathsf{inf}(M' \cup M'') \geqslant \mathsf{inf}\, M$ we have that $\mathsf{inf}(M' \cup M'') \neq \overleftarrow{m}$ and, thus, there is a corner of $\vec{x}$ and $\vec{y}$ in $\vec{S} \setminus \{ \vec{m}, \overleftarrow{m} \}$.

It remains the case that $|M| = 1$, say $M = \{\vec{m}\}$. In this case however, we have that $\vec{s} \leqslant \vec{m}$ for every $\vec{s} \in \vec{S}$. If $\vec{S} = \{ \vec{m}, \overleftarrow{m} \}$ the statement is trivial, so let $\vec{s} \in \vec{S} - \vec{m}$ be $\leqslant$-maximal such that $\vec{s} \neq \overleftarrow{m}$. Such an $\vec{s}$ exists as $\overleftarrow{m}$ is a $\leqslant$-minimal element of $\vec{S}$. Then $\vec{m}$ is the unique upper-cover of $\vec{s}$. Thus $\vec{S} - \vec{s}$ is a woven poset by Lemma 9.16. Moreover, $\overleftarrow{m}$ is the unique lower cover of $\overleftarrow{s}$ and, since $\overleftarrow{m} \neq \vec{s}$ it is also the unique lower cover of $\overleftarrow{s}$ in $\vec{S} - \vec{s}$. Thus, $(\vec{S} - \vec{s}) - \overleftarrow{s}$ is a woven poset by Lemma 9.16, and $S - s$ is a submodular separation system. $\square$

The Dedekind-MacNeille completion of posets [16], which we have seen in Section 8.2, allows us to embed every woven poset into a lattice so that the poset is woven in this lattice. We showed in Section 8.2 that this completion can also be applied to submodular separation systems to obtain a universe of separations in which the separation system is submodular.

In particular, if P is a woven poset and $p \in P$ such that $P' = P - p$ is again woven, there are lattices L and L' such that P is woven in L and P' is woven in L'. If we could arrange for these two lattices to be sublattices of one another, $L' \subseteq L$, in such a way that every element of $P' \subseteq L'$ is mapped to the corresponding element of $P \subseteq L$, then this would imply that P could be unravelled as a woven subset of L in the sense of Problem 9.2.

The way in which we constructed P', however, makes this almost impossible. We choose p as an element with a unique upper, or a unique lower cover. Now if $p \in P$ has a unique upper cover q, say, and is also the supremum of some two points $r, s \in P \setminus \{p\}$, then the Dedekind-MacNeille completion L' of P' cannot be embedded in the way outlined above into the Dedekind-MacNeille completion L of P: in L', the images of r and s have the image of q as supremum and an embedding as a sublattice would need to preserve this property, but the images of r and s in L have the image of p as their supremum. (However, L' *is* order-isomorphic to a sub*poset* of L.)

Bibliography

[1] S. Arnborg, D. G. Corneil, and A. Proskurowski. (1987). *Complexity of finding embeddings in a k-tree*, SIAM Journal on Algebraic Discrete Methods **8**, no. 2, 277–284, doi:10.1137/0608024.

[2] A. Berry, R. Pogorelcnik, and G. Simonet. (2010). *An introduction to clique minimal separator decomposition*, Algorithms **3**, no. 2, 197–215, doi:10.3390/a3020197.

[3] T. S. Blyth and M. F. Janowitz. (1972). *Residuation theory*, International Series of Monographs on Pure and Applied Mathematics, vol. 102, Pergamon.

[4] N. Bowler and J. Kneip. (2020). *Separations of sets*, Order **37**, no. 2, 411–425, doi:10.1007/s11083-019-09512-4.

[5] C. Bürger and J. Kurkofka. (2020). *Duality theorems for stars and combs III: Undominated combs*, preprint at arXiv:2004.00592.

[6] ———. (2021). *Duality theorems for stars and combs I: Arbitrary stars and combs*, Journal of Graph Theory, 1–30, doi:10.1002/jgt.22706.

[7] ———. (to appear). *Duality theorems for stars and combs II: Dominating stars and dominated combs*, Journal of Graph Theory, preprint at arXiv:2004.00593.

[8] ———. (to appear). *Duality theorems for stars and combs IV: Undominating stars*, Journal of Graph Theory, preprint at arXiv:2004.00591.

[9] J. Carmesin. (2016). *A short proof that every finite graph has a tree-decomposition displaying its tangles*, European Journal of Combinatorics **58**, 61–65, doi:10.1016/j.ejc.2016.04.007.

[10] ———. (2019). *All graphs have tree-decompositions displaying their topological ends*, Combinatorica **39**, no. 3, 545–596, doi:10.1007/s00493-018-3572-0.

[11] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. (2016). *Canonical tree-decompositions of finite graphs I. Existence and algorithms*, Journal of Combinatorial Theory, Series B **116**, 1–24, doi:10.1016/j.jctb.2014.04.001.

[12] ———. (2016). *Canonical tree-decompositions of finite graphs II. Essential parts*, Journal of Combinatorial Theory, Series B **118**, 268–283, doi:10.1016/j.jctb.2014.12.009.

[13] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. (2014). *Connectivity and tree structure in finite graphs*, Combinatorica **34**, no. 1, 11–46, doi:10.1007/s00493-014-2898-5.

[14] J. Carmesin, M. Hamann, and B. Miraftab. (2020). *Canonical trees of tree-decompositions*, preprint at arXiv:2002.12030.

[15] D. Coudert and G. Ducoffe. (2016). *Clique-decomposition revisited*, INRIA Sophia Antipolis - I3S. Available at hal-01266147.

[16] B. A. Davey and H. A. Priestley. (2002). *Introduction to lattices and order*, 2nd ed., Cambridge University Press.

[17] W. Dicks and M. J. Dunwoody. (1989). *Groups acting on graphs*, Cambridge Studies in Advanced Mathematics, Cambridge University Press.

[18] R. Diestel. (2017). *Ends and tangles*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **87**, no. 2, 223–244, doi:10.1007/s12188-016-0163-0.

[19] ———. (2017). *Graph theory*, 5th ed., Graduate Texts in Mathematics, vol. 173, Springer, Berlin, Heidelberg.

[20] ———. (2018). *Abstract separation systems*, Order **35**, 157–170, doi:10.1007/s11083-017-9424-5.

[21] ———. (2018). *Tree sets*, Order **35**, no. 1, 171–192, doi:10.1007/s11083-017-9425-4.

[22] ———. (2019). *Tangles in the social sciences*, preprint at arXiv:1907.07341.

[23] R. Diestel, P. Eberenz, and J. Erde. (2017). *Duality theorems for blocks and tangles in graphs*, SIAM Journal on Discrete Mathematics **31**, no. 3, 1514–1528, doi:10.1137/16M1077763.

[24] R. Diestel, J. Erde, C. Elbracht, and M. Teegen. *Duality and tangles of set partitions*, in preparation.

[25] R. Diestel, J. Erde, and D. Weißauer. (2019). *Structural submodularity and tangles in abstract separation systems*, Journal of Combinatorial Theory, Series A **167C**, 155–180, doi:10.1016/j.jcta.2019.05.001.

[26] R. Diestel, F. Hundertmark, and S. Lemanczyk. (2019). *Profiles of separations: in graphs, matroids, and beyond*, Combinatorica **39**, no. 1, 37–75, doi:10.1007/s00493-017-3595-y.

[27] R. Diestel and S. Oum. (2019). *Tangle-tree duality in graphs, matroids and beyond*, Combinatorica **39**, no. 4, 879–910, doi:10.1007/s00493-019-3798-5.

[28] ———. (2021). *Tangle-tree duality in abstract separation systems*, Advances in Mathematics **377**, 107470, doi:10.1016/j.aim.2020.107470.

[29] R. Diestel and G. Whittle. (2016). *Tangles and the Mona Lisa*, preprint at arXiv:1603.06652.

[30] M. J. Dunwoody and B. Krön. (2015). *Vertex cuts*, Journal of Graph Theory **80**, no. 2, 136–171, doi:10.1002/jgt.21844.

[31] P. Eberenz. (2015). *Characteristics of profiles*, MSc thesis, Universität Hamburg. Retrieved from `https://www.math.uni-hamburg.de/spag/dm/papers/Eberenz.pdf` on April 20, 2021.

[32] C. Elbracht. (2017). *Tangles determined by majority vote*, MSc thesis, Universität Hamburg. Retrieved from `https://www.math.uni-hamburg.de/spag/dm/papers/Elbracht_MSc.pdf` on April 20, 2021.

[33] C. Elbracht, D. Fioravanti, S. Klepper, J. Kneip, L. Rendsburg, M. Teegen, and U. v. Luxburg. (2020). *Tangles: From weak to strong clustering*, preprint at arXiv:2006.14444.

[34] C. Elbracht and J. Kneip. (to appear). *A canonical tree-of-tangles theorem for submodular separation systems*, Combinatorial Theory, preprint at arXiv:2009.02091.

[35] C. Elbracht, J. Kneip, and M. Teegen. (2020). *A note on generic tangle algorithms*, preprint at arXiv:2005.14648.

[36] ———. (2020). *Tangles are decided by weighted vertex sets*, Advances in Combinatorics **9**, doi:10.19086/aic.13691.

[37] ———. (2021). *The structure of submodular separation systems*, preprint at arXiv:2103.13162.

[38] ———. (2021). *Trees of tangles in abstract separation systems*, Journal of Combinatorial Theory, Series A **180**, 105425, doi:10.1016/j.jcta.2021.105425.

[39] ———. (2021). *The unravelling problem*, preprint at arXiv:2103.12565.

[40] ———. (to appear). *Obtaining trees of tangles from tangle-tree duality*, Journal of Combinatorics, preprint at arXiv:2011.09758.

[41] ———. (to appear). *Trees of tangles in infinite separation systems*, Mathematical Proceedings of the Cambridge Philosophical Society, preprint at arXiv:2005.12122.

[42] C. Elbracht, J. Kurkofka, and M. Teegen. (2021). *Edge-connectivity and tree-structure in finite and infinite graphs*, preprint at arXiv:2012.07651.

[43] A.-K. Elm and J. Kurkofka. (2020). *A tree-of-tangles theorem for infinite tangles*, preprint at arXiv:2003.02535.

[44] E. Fluck. (2019). *Tangles and single linkage hierarchical clustering*, 44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019), pp. 38:1–38:12, doi:10.4230/LIPIcs.MFCS.2019.38.

[45] J. Geelen. (2016). *Personal communication*, Oberwolfach.

[46] J. Geelen, B. Gerards, N. Robertson, and G. Whittle. (2006). *Obstructions to branch-decomposition of matroids*, Journal of Combinatorial Theory, Series B **96**, no. 4, 560–570, doi:10.1016/j.jctb.2005.11.001.

[47] J. Geelen, B. Gerards, and G. Whittle. (2009). *Tangles, tree-decompositions and grids in matroids*, Journal of Combinatorial Theory, Series B **99**, no. 4, 657–667, doi:10.1016/j.jctb.2007.10.008.

[48] A. C. Giannopoulou, K. Kawarabayashi, S. Kreutzer, and O. Kwon. (2020). *The canonical directed tree decomposition and its applications to the directed disjoint paths problem*, preprint at arXiv:2009.13184.

[49] J. P. Gollin and J. Kneip. (2021). *Representations of infinite tree sets*, Order **38**, no. 1, 79–96, doi:10.1007/s11083-020-09529-0.

[50] M. Grohe. (2016). *Tangled up in blue (A survey on connectivity, decompositions, and tangles)*, preprint at arXiv:1605.06704.

[51] M. Grohe and P. Schweitzer. (2016). *Computing with tangles*, SIAM Journal on Discrete Mathematics **30**, no. 2, 1213–1247, doi:10.1137/15M1027565.

[52] G. A. Grätzer. (1978). *General lattice theory*, Pure and applied mathematics: a series of monographs and textbooks, Academic Press, New York.

[53] R. Halin. (1964). *Über simpliziale Zerfällungen beliebiger (endlicher oder unendlicher) Graphen*, Mathematische Annalen **156**, no. 3, 216–225, doi:10.1007/BF01363288 (German).

[54] ______. (1991). *Lattices of cuts in graphs*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **61**, 217–230, doi:10.1007/BF02950767.

[55] ______. (1991). *Tree-partitions of infinite graphs*, Discrete Mathematics **97**, no. 1-3, 203–217, doi:10.1016/0012-365X(91)90436-6.

[56] T. Johnson, N. Robertson, P. D. Seymour, and R. Thomas. (2001). *Directed tree-width*, Journal of Combinatorial Theory, Series B **82**, no. 1, 138–154, doi:10.1006/jctb.2000.2031.

[57] J. Kneip. (2019). *Ends as tangles*, preprint at arXiv:1909.12628.

[58] ______. (2020). *Tangles and where to find them*, Doctoral thesis (Dissertation), Universität Hamburg, Staats- und Universitätsbibliothek Hamburg Carl von Ossietzky. urn:nbn:de:gbv:18-ediss-89005.

[59] F. Lazebnik and V. A. Ustimenko. (1995). *Explicit construction of graphs with an arbitrary large girth and of large size*, Discrete Applied Mathematics **60**, no. 1-3, 275–284, doi:10.1016/0166-218X(94)00058-L.

[60] H.-G. Leimer. (1993). *Optimal decomposition by clique separators*, Discrete Mathematics **113**, no. 1-3, 99–123, doi:10.1016/0012-365X(93)90510-Z.

[61] C.-H. Liu. (2015). *Packing and covering immersions in 4-edge-connected graphs*, preprint at arXiv:1505.00867.

[62] H. M. MacNeille. (1937). *Partially ordered sets*, Transactions of the American Mathematical Society **42**, no. 3, 416–460, doi:10.2307/1989739.

[63] B. Reed. (1999). *Introducing directed tree width*, Electronic Notes in Discrete Mathematics **3**, 222–229, doi:10.1016/S1571-0653(05)80061-7. 6th Twente Workshop on Graphs and Combinatorial Optimization.

[64] N. Robertson and P. D. Seymour. (1991). *Graph minors. X. Obstructions to tree-decomposition*, Journal of Combinatorial Theory, Series B **52**, no. 2, 153–190, doi:10.1016/0095-8956(91)90061-N.

[65] D. Seese. (1985). *Tree-partite graphs and the complexity of algorithms*, Fundamentals of Computation Theory, pp. 412–421, doi:10.1007/BFb0028825.

[66] R. E. Tarjan. (1985). *Decomposition by clique separators*, Discrete Mathematics **55**, no. 2, 221–232, doi:10.1016/0012-365X(85)90051-2.

[67] A. W. Tucker. (1957). *Dual systems of homogeneous linear relations*, Linear inequalities and related systems. (AM-38), pp. 1–18, doi:10.1515/9781400881987-002.

[68] K. Wagner. (1937). *Über eine Eigenschaft der ebenen Komplexe*, Mathematische Annalen **114**, 570–590, doi:10.1007/BF01594196 (German).

[69] P. Wollan. (2015). *The structure of graphs not admitting a fixed immersion*, Journal of Combinatorial Theory, Series B **110**, 47–66, doi:10.1016/j.jctb.2014.07.003.

A. Appendix

A.1. Deutschsprachige Zusammenfassung

Diese Dissertation beschäftigt sich mit *Tangles* in abstrakten Teilungssystemen, einer Charakterisierung von dichten, oder besonders zusammenhängenden, Teilstrukturen – also „Clustern“ – für diverse Kontexte, sowie mit den Submodularitätsanforderungen die man an diese Teilungssysteme stellen kann. Tangles sind eine indirekte Art Cluster zu beschreiben: Sie geben nicht explizite Teile einer Struktur an, sondern machen nur für verschiedene Weisen eine Struktur zu Zerteilen eine Angabe in welchem Teil der Cluster liegt. Dieses Vorgehen erlaubt es auch Cluster zu erfassen, bei denen die konkrete Zugehörigkeit einzelner Punkte ungewiss ist. Tangles stammen ursprünglich aus der Graphen-Minorentheorie, wo sie genutzt werden um die Baumartigkeit eines Graphen zu untersuchen. Insgesamt ist die vorliegende Arbeit in drei Teile gegliedert.

In Teil I geht es – in zwei unteschiedlichen Kontexten – um Strukturen die ein einzelnes Tangle induzieren kann.

Eine offene Frage über Graphen-Tangles ist jene nach der Existenz einer *Entscheidermenge*, einer Menge X von Ecken, so dass, wannimmer das Tangle eine Teilung (A, B) nach B orientiert, die Mehrheit von X auf der B-Seite liegt. In Kapitel 3 beweisen wir die Existenz einer gewichteten Version von Entscheidern, also einer nicht-negativen Funktion auf den Ecken so dass die Summe der Funktionswerte über A bei einer Teilung (A, B) im Tangle stets kleiner als die entsprechende Summe über B ist.

Kapitel 4 behandelt Dualität zwischen Tangles auf den zwei Seiten eines bipartiten Graphen. Für einen bipartiten Graphen mit Partitionsklassen X und Y zeigen wir wie man Tangles der Teilungen von X bzw. Y so definieren kann, dass auf ganz natürliche Weise die Tangles auf X welche auf Y induzieren und umgekehrt. Es stellt sich dann in Abschnitt 4.2 heraus, dass diese dualen Tangles sich auch durch eine neue Art von Tangle, welches auf Teilungen der Kanten definiert ist, bezeugen lassen.

Teil II ist den *Trees of Tangles* gewidmet. Der Tree-of-Tangles-Satz ist einer der fundamentalen Sätze der Tangletheorie. Man versteht unter diesem Begriff tatsächlich eine ganze Familie von Sätzen, wovon jeder, auf die eine oder andere Weise, aussagt, dass es zu jeder (wohlgeformten) Menge $\mathcal{T}$ von Tangles eine Menge N verschachtelter Teilungen gibt, so dass je zwei Tangles aus $\mathcal{T}$ sich durch eine Teilung in N unterscheiden. Die verschachtelte Anordnung der Teilungen offenbart dabei eine baumförmige Anordnung der Tangles in $\mathcal{T}$. Je nachdem wie wohlgeformt die Menge $\mathcal{T}$ ist, liefern diese Sätze zusätzliche Eigenschaften: dass N kanonisch ist und/oder dass dir Unterscheiderteilungen effizient (d.h. besonders schlank) sind.

In Kapitel 5 führen wir mit den *Splinter-Lemmata* ein allgemeines Prinzip ein, nach dem wir die bedeutsamsten dieser Sätze einheitlich und einfach beweisen können. Durch diese Lemmata können wir auch schnell neue Tree-of-Tangles-Sätze schaffen, beispielsweise einen für Profile in aufsteigenden Folgen submodularer Teilungssysteme.

Dieses Framework erweitern wir in Abschnitt 5.8 und 5.9 auch auf unendliche Teilungssysteme und Tanglemengen. Mit dieser Erweiterung beweisen wir zwei Existenzsätze für bestimmte Baumzerlegungen unendlicher Graphen neu, indem wir einen neuen Satz entwickeln der in Mitte zwischen den beiden liegt. Eine weitere Anwendung dieses Splinter-Lemmas fürs Unendliche, nämlich auf Kantenzusammenhang in unendlichen Graphen, stellen wir in Abschnitt 5.10 vor.

Auch Kapitel 6 besteht aus Beweisen von Tree-of-Tangles-Sätzen. Die Sätze, die wir dort zeigen, sind für sich genommen nichts besonderes – gerade vor dem Hintergrund von Kapitel 5. Die Art und Weise auf die wir sie beweisen ist jedoch außergewöhnlich: Wir wenden den zweiten fundamentalen Satz der Tangletheorie an, den Tangle–Tree-Dualitätssatz. Zwar können beide Sätze, der Tree-of-Tangles-Satz und der Dualitätssatz, einen Baum als Ergebnis liefern, aber gewöhnlich tut der Dualitätssatz dies nur, wenn es keinen Tangle gibt, während der Tree-of-Tangles-Satz nur Anwendung findet, wenn es mehrere Tangles gibt. In Kapitel 6 reizen wir den Tangle–Tree-Dualitätssatz an seine Grenzen aus und bringen ihn dazu einen Tree of Tangles zu liefern.

Abschließend stellen wir in Kapitel 7 einen heuristischen Ansatz zur Berechnung von Tangles und Trees of Tangles vor. Wir erklären, wie man ohne perfektes Kentnis eines vollständigen Teilungssystems dennoch eine Annäherung der Tangles und Trees of Tangles für diese berechnen kann.

Teil III befasst sich mit den Teilungssystemen an sich – genauer, mit ihrer *Submodularität*. Submodularität ist eine wichtige Eigenschaft für die Tangletheorie und spielt in nahezu jedem Beweis über Tangles eine Rolle. Es gibt jedoch verschiedene Abstufungen von Submodularitätsbedingungen, die man an sein Teilungssystem stellen kann und Kapitel 8 untersucht die Beziehung zwischen den drei natürlichen Submodularitätsbedingungen: *Submodularität*, *(strukturelle) Submodularität innerhalb eines Teilungsuniversums* und *ordnungsinduzierte Submodularität* – eine jede Bedingung schärfer als die vorhergehenden. Insbesondere zeigen wir in Abschnitt 8.2, wie die (für sich genommene) Submodularität eines Teilungssystem auch durch die Konstruktion eines geeigneten Teilungsuniversums um das Teilungssystem herum bezeugt werden kann und verknüpfen dadurch die schwächste mit der mittleren Bedingung. Andererseits beweisen wir in Abschnitt 8.3, dass ordnungsinduzierte Submodularität eine echt stärkere Eigenschaft ist als bloße Submodularität innerhalb eines Teilungsuniversums, was die stärkste Bedingung von der mittleren abgrenzt. Über diesen Vergleich der Bedingungen hinaus entwicklen wir zwei Zerlegungssätze für Teilungssysteme die submodular im zweiten Sinne in einem distributiven Teilungsuniversum sind.

In Kapitel 9 stellen wir eine weiter fundamentale Frage zu submodularen Teilungssystemen: Gibt es in einem submodularen Teilungssystem immer eine Teilung, die wir löschen können, so dass dass der Rest noch submodular ist? Wenn wir eine Folge solcher Teilungen finden können, die wir nacheinander unter Beibehaltung der Submodularität löschen können bis wir schlussendlich die leere Menge erreichen, so sprechen wir von einer *Entwirrung*. Wir beweisen die Existenz solcher Entwirrungsfolgen, sowohl für das stärkste, als auch das schwächste Submodularitätskonzept, in den Abschnitten 9.2 und 9.4. Für das dritte Konzept von Submodularität, Submodularität innerhalb eines Teilungsuniversums, präsentieren wir ein Gegenbeispiel, das keine Entwirrung ermöglicht.

A.2. Summary

This thesis treats *tangles* in abstract separation systems, a way of characterising dense or well-connected substructure or 'clusters' in various contexts, and the underlying submodularity properties that one can demand of these separation systems. Tangles describe clusters in an indirect way: not by specifying explicit parts of a data structure, but only by declaring for every way of cutting the structure into two parts which side the cluster lies. This approach allows to capture concepts of clusters where the membership of individual points in a specific cluster is uncertain. The concept of tangles originates from graph minor theory, where they are used to characterise how tree-like a graph is. This thesis is divided into three parts.

In Part I we present two instances of structure induced by a single tangle.

An open question about tangles in graphs is whether each of them has a *decider set*: a set X of vertices so that for every separation (A, B) where the tangle points to B more vertices of X are in B than in A. In Chapter 3 we show that a weighted version of deciders exists, i.e., that there always is a non-negative weight function defined on the vertices such that, for every separation (A, B) in the tangle, the sum over the weights in B is larger than the sum over the weights in A.

Chapter 4 covers dual tangles in the setting of bipartite graphs. Given a bipartite graph with partition classes X and Y, we show how one can define tangles in the separations of X and Y, so that tangles in X induce tangles in Y – and vice versa – in a natural way. These dual tangles turn out to also be witnessed by a new kind of tangles defined on the set separations of the edge set, which we cover in Section 4.2.

Part II, is devoted to trees of tangles. The tree-of-tangles theorem is one of the fundamental theorems of tangle theory. In fact, the term is used to describe a whole class of theorems, all of which state that for every (sufficiently nice) set of tangles $\mathcal{T}$, then there exists a nested set of separations N containing a distinguishing separation for every pair of tangles in $\mathcal{T}$. Such a nested set reveals a tree-shaped arrangement of the tangles in the separation system. Depending on the conditions on the set of tangles $\mathcal{T}$, these theorems guarantee additional properties of the set N: canonicity and/or the efficiency of the distinguishing separations.

Chapter 5 introduces our *splinter lemmas*, a framework which allows us to prove in a unified and simple way the most relevant existing tree-of-tangles theorems. They also allow as introduce some of our own tree-of-tangles theorems; for example, for profiles in sequences of separation systems.

We also extend this framework to an infinite setting in Sections 5.8 and 5.9, re-proving two existence theorems on tree-decompositions of infinite graphs by introducing a new theorem which sits midway between the two. We introduce another application, this time to edge-connectivity in infinite graphs, in Section 5.10.

Chapter 6 comprises proofs of tree-of-tangles theorems as well. And while the theorems themselves are unspectacular, especially in light of Chapter 5, the extraordinary aspect is the method by which we prove those theorems: by applying the tangle–tree duality theorem. The tangle–tree duality theorem and the tree-of-tangles theorem are the two fundamental pillars on which tangle theory stands. But while both can give a tree-like arrangement of separations as a result, the way in which you apply them is very different: the tree-of-tangles theorem gives you tree if there are tangles, the tangle–tree duality theorem gives you a tree if there is no tangle. In Chapter 6 we push the limits of the tangle–tree duality theorem and make it produce a tree of tangles.

Finally, Chapter 7 introduces a heuristic approach for computing tangles and trees of tangles. There we explain how, without perfect knowledge of a complete separation system, one can compute an approximation of the tangles and build trees of tangles for such approximate tangles.

The final part, Part III, is concerned with the structural properties of the separation systems – specifically with *submodularity*. Submodularity is a central property in tangle theory and features in virtually every proof concerned with tangles. However, there are multiple gradations of submodularity conditions. Chapter 8 explores the relationship between the three natural submodularity conditions that one can impose onto a separation system: *submodularity*, *(structural) submodularity inside a universe of separations*, and *order-induced submodularity*, where each is stronger requirement than the previous. In particular, in Section 8.2 we show how the submodularity of a separation system on its own can be witnessed by constructing an appropriate universe around it, linking the first and second condition. On the other hand, we show in Section 8.3 that order-induced submodularity is a strictly stronger property than mere submodularity in a universe, making the distinction between the second and third condition explicit. Beyond comparing the types of submodularity, we develop two decomposition theorems for separation systems which satisfy submodularity of the second type within a distributive universe of separations.

Chapter 9 introduces one fundamental question about submodular separation systems: Given a submodular separation system, is there always a separation that we can delete, so that the remainder is still submodular? If we can find a sequence of separations that we can delete one after the other while maintaining submodularity until we reach the empty set, we speak of an *unravelling*. We show the existence of unravellings for the strongest and for the weakest concept of submodularity in Section 9.2 and Section 9.4, respectively. For the third concept, (structural) submodularity inside a universe of separations, we give a counterexample in Section 9.3.

A.3. Publications related to this thesis

The following articles and preprints are related to this dissertation:

Chapter 3 is based on [36].
Chapter 4 is based on [24].
Chapter 5 up to and including Section 5.5 is based on [38].
Sections 5.8 and 5.9 are based on [41].
Section 5.10 is based on [42].
Chapter 7 is partially based on [35] as elaborated in the following section.
Chapter 8 is based on [37].
Chapter 9 is based on [39].

A.4. Declaration of contributions

Most research presented in this thesis was conducted in close collaboration with my respective co-authors of the publications listed in Appendix A.3. This is especially true for Christian Elbracht and Jakob Kneip, both of whom I head the delight of sharing an office with. Thus, we often contributed equal amounts of work in research and writing. Even if one of us had the final idea which made a proof work, this was often a product of many previous failed iterations contributed by the others. I will now go into more detail on the work that went into the respective chapters of this thesis, point out where contributions differ, and emphasise some of the ideas I contributed myself.

Chapter 3

Chapter 3 was joint work with Christian Elbracht and Jakob Kneip. We had worked on the problem of deciders for about a year in close collaboration and finally achieved a breakthrough on the day that I suggested trying for a linear programming approach.

Our first proof sketch was rather complicated, Elbracht then found Tucker's Lemma which completed and simplified the proof to its current form. Kneip touched up our notes from that common research into what is now all of Chapter 3 excepting Section 3.2. Section 3.2, which is a later addition, was drafted by Elbracht, the final version is written by me.

Chapter 4

Chapter 4 is based on a section of [24], a joint paper of Reinhard Diestel, Christian Elbracht, Joshua Erde, and me. Note that I suggested both Diestel and Erde as evaluators for this thesis, and that Diestel is my doctoral supervisor. The general concept of dual separation systems was introduced by Diestel, and in that paper he introduces his work on them. The results presented in this thesis is based on joint research with only Erde and Elbracht which is largely independent of Diestel's work on dual separation system apart from its motivation; the remaining parts of the paper [24] are written by Diestel.

Erde introduced Elbracht and me to the concept of tangles on the sides of a bipartite graph and showed us a proof that one could shift a tangle to give a tangle again, the predecessor of what is now Theorem 5 and asked us whether shifting stabilizes. Elbracht and I then proved Theorem 6; we also discovered and analysed the associated edge tangles, which is Section 4.2. The part of [24] that is presented in Chapter 4 is based on a draft by Elbracht and me and was finished up jointly by Elbracht, Erde, and myself.

Chapter 5 up to and including Section 5.5

The splinter lemma and the canonical splinter lemma are joint work with Elbracht and Kneip [38]. The first version of Lemma 13 was discovered by Elbracht and myself in close cooperation. It was first formulated in terms of profiles and with a more complex proof along the lines of Lemma 5.5. After presenting the proof to Kneip, he greatly simplified both the statement and the proof to the form they have today.

The canonical version was developed, again in close cooperation, of Elbracht and me, although our work on canonicity in general was joint among all three of us. Elbracht and I prepared a first draft of [38], the finished version was written mostly by Kneip.

Sections 5.6 and 5.7

The relation splinter lemma was developed by Elbracht, Kneip, and me during the writing of [38]. The application to directed tangles is a recent addition, based on joint and equal work of Elbracht and me. Section 5.7 is based on shared notes of ours, the final presentation in this thesis is my own.

Sections 5.8 and 5.9

Section 5.8 is based on [41] which is joint work with Elbracht and Kneip. Kneip developed the splinter lemma for profinite universes, which is not part of this thesis. The part presented here was developed by Elbracht and me. I developed a topological approach which was later superseded by Elbrachts definition of 'thin splintering'. Elbracht also devised the idea of applying the lemma to separators instead of separations. The parts of [41] that are contained in this thesis are written mostly by Elbracht and me in close collaboration, based on drafts written by Elbracht.

Section 5.10

Section 5.10 is joint work with Jan Kurkofka and Christian Elbracht. Kurkofka asked Elbracht and me for a tree-of-tangles theorem for edge blocks, upon which Elbracht and me drafted the proof of Theorem 25. The relation to tree-cut decompositions was observed by Kurkofka, who also wrote most of the final version. The three of us then proved Theorem 26 in close cooperation.

Chapter 6

Chapter 6 is joint work with Elbracht and Kneip based on a question originally posed by Nathan Bowler and Joshua Erde of whether it is possible to obtain a tree-of-tangles theorem from clever application of the tangle–tree duality theorem. Elbracht and me jointly developed such a method, which is the first proof of Theorem 6.10 in Section 6.2. Kneip then devised the proof based on uncrossing stars and derived the bound on degrees in trees of tangles as in Section 6.3. He also wrote most of the corresponding sections. The generalization to efficient distinguishers (Sections 6.4 and 6.5) as well as the mixed order case (Section 6.6) were created solely by Elbracht and me, and we wrote Sections 6.4 to 6.6 in close collaboration.

Chapter 7

The work on algorithms for Chapter 7 was together with Elbracht and Kneip and primarily consisted of joint programming. During our work we developed the simple tangle search algorithm (the one presented in Section 7.2) as well as a more complicated tree-of-tangles algorithm; the latter was developed by myself with some important contributions from Elbracht. Elbracht and I wrote a draft [35] together presenting those results. Section 7.3 with its on-line algorithm and Section 7.4 with its simplified tree-of-tangles algorithm are my own work.

Chapter 8

Chapter 8 is based on joint research with Elbracht and Kneip. Section 8.2 is by Kneip and me. The original question and the idea to apply the Dedekind–MacNeille-completion was devised by me. We then developed the proof of Theorem 8.2 together.

Sections 8.3 to 8.5 are by only Elbracht and me. The concept of submodular decompositions in Section 8.5 as well as an earlier version of Theorem 8.11 was proven by me. Upon me mentioning the result to him, Elbracht promptly proved Theorem 35, which allowed me to strengthen Theorem 8.11 to be a decomposition into three (instead of six) parts, and lead me to discover Theorem 34.

Section 8.6 is my own work.

Chapter 9

Chapter 9 is also based on joint work with Elbracht and Kneip. The question of deletable separations was posed by Joshua Erde. Christian Elbracht and me discovered Lemma 9.7 and Proposition 9.3 together. Elbracht, Kneip, and I then further pursued the issue together and Kneip proved the remaining results of Section 9.2 (originally for universes instead of lattices) by proving an earlier version of Theorem 9.8, while the research for Sections 9.3 and 9.4 was done by only Elbracht and myself.

A.5. Acknowledgements

I'd like to thank my supervisor Reinhard Diestel for his advice and support throughout my studies. Thank you for introducing me to the joys of graph theory!

I'd like to thank my collaborators for our enjoyable and exciting work together. In particular, I would like to thank Christian and Jakob, my office-mates, for the excellent teamwork and delightful conversations – be they about mathematics or otherwise. You became good friends to me.

Indeed, I would also like to thank all my wonderful friends – those that I've met during university and those I had made before. Thank you all for being there, enjoying the after-work hours, going out, cooking, and playing games together with me. Over this past year of isolation I have missed all of this very dearly and I hope to see you again soon.

I would like to thank my colleagues for creating a pleasant working environment, for always being ready to hear my thoughts on a mathematical problem and giving their own thoughts, for inspiring and fun conversations over lunch and coffee.

Lastly, I'd like to thank my family. Thank you for your love, for you care, and for your support! I send special thanks in faith they reach my late mother, in loving memory.

A.6. Eidesstattliche Erklärung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Dissertation selbst verfasst und keine anderen als die angegebenen Hilfsmittel benutzt habe. Darüber hinaus versichere ich, dass diese Dissertation nicht in einem früheren Promotionsverfahren eingereicht wurde.

Maximilian Teegen