Distinguishing and witnessing dense structures in graphs and abstract separation systems

Dissertation

zur Erlangung der Würde des
Doktors der Naturwissenschaften

der Fakultät für Mathematik, Informatik
und Naturwissenschaften,

Fachbereich Mathematik der

Universität Hamburg

vorgelegt von

Christian Elbracht

aus Hamburg

Hamburg
2021

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

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

Zweitgutachter:
Joshua Erde, PhD (Technische Universität Graz)

Datum der Disputation:
02.09.2021

Contents

Chapter 1

Introduction

1.1 Historic overview

One of the most fundamental areas of study in graph theory is the analysis of a graph's connectivity. Most results in that area belong, roughly speaking, to one of two types: the first type is concerned with finding a highly connected substructure of the graph and wants, in the case of a graph not containing such a structure, to find a 'dual structure' which witnesses the absence of that highly connected substructure. Perhaps the first result of this type is Menger's theorem which characterizes high connectivity between two vertices via the existence of a lot of disjoint paths. It states (in one form) that, for any $k \in \mathbb{N}$ and any two vertices of a graph, there always exist either k disjoint paths between the two vertices or a separator of size less than k which separates the two vertices from another. Thus, we either find a structure witnessing the high connectivity between the two vertices (the k disjoint paths) or we find a dual structure which witnesses that such paths cannot exist (the separator of size $< k$).

The other type of connectivity result is not interested in finding a highly connected substructure, but instead wants to decompose the graph into small parts, so that each part contains just one such substructure. As a trivial example, every graph is the disjoint union of its components, so we can 'decompose' the graph into its 1-connected pieces. The well-known block-cutvertex-tree extends this, in a way, to 2-connected pieces. It, roughly speaking, allows one to decompose, in a tree-like way, a connected graph along its cut vertices into 2-connected components.

These two primary results of the second type motivated, starting in the 1960s, the search for corresponding analogues for higher connectivity: we would like to be able to decompose a graph in a tree-like way into its $(k+1)$-connected pieces. However, for increasing k, this task becomes more and more challenging if one interprets '$(k+1)$-connected pieces' as '$(k+1)$-connected subgraphs'. Already for $k = 2$, this does no longer seem to be possible in this naive way. Thus, the well-known decomposition theorem by Tutte [74] does not give a decomposition into 3-connected subgraphs: instead Tutte showed (formulated in modern terms) that every graph admits a tree-decomposition of adhesion 2 such that every torso of that decomposition is either 3-connected or a cycle.

Moving from 3-connected subgraphs to the torsos in this theorem is a first step in an important shift in perspective: instead of requiring that the parts of the decomposition are 3-connected 'in their own', Tutte's theorem only guarantees that they are 3-connected 'in the surrounding graph'.

This shift in perspective resulted in considering k-inseparable sets in a graph instead of k-connected subgraphs. The first one to consider such sets was Mader [65]. Formally, a *k-inseparable* set X in a graph G is a vertex set of size at least $k+1$ with the property that no two vertices in X can be separated in G by k vertices. Nowadays, maximal such sets are called *blocks*. Dunwoody and Kröhn [33] noticed that these k-blocks can be used to generalize, in a way, the classic result from Tutte to greater values of k, by showing that the k-blocks of a finite k-connected graph can be separated from another in a tree-like way. This result has then been generalized by Carmesin, Diestel, Hundertmark and Stein [13] to graphs of arbitrary connectivity.

Independently of this theory of k-blocks, Robertson and Seymour [68] invented a new, and radically different notion of a 'k-connected piece' as a tool in their graph minor project: *tangles*. Instead of characterizing a highly connected substructure explicitly by some fixed set of vertices and/or edges, they capture a k-connected piece in an indirect way, following the idea that, for every small (that is of size less than k) vertex set X of our graph, the majority of a k-connected piece should be contained in one component of $G - X$. This can easily be modelled via *separations* of a graph: a separation (A, B) of order k in a graph G consists of two vertex sets A and B such that $A \cup B$ is the whole vertex set of G, the intersection $A \cap B$ has size k, and there are no edges from $A \setminus B$ to $B \setminus A$. Following the idea that a k-connected piece of a graph should always be contained, with its majority, in one of the two sides, A or B, of such a separation, Robertson and Seymour then defined their notion of a 'k-connected piece' via the separations of order less than k: a *k-tangle* is a way to choose a side for each such separation in a 'consistent way', which is traditionally given by requiring that a tangle should not contain any three separations such that their small sides, where the small side of a separation (A, B) in a tangle is $G[A]$, cover the whole graph.

This novel way of indirectly capturing the highly connected pieces of a graph in particular includes another prototypical example of a substructure of a graph, of which we might think of as 'highly connected', but which is not captured by the notion of a k-block: large grids. Every large enough grid induces a k-tangle by orienting every separation towards the side containing more of that grid, however a large grid does not contain any k-block for $k > 4$. In fact, the famous grid theorem by Robertson and Seymour implies that the existence of a high order tangle is quantitatively bound to the existence of a large grid minor: a graph contains a large grid as a minor if and only if the graph contains a high order tangle.

Robertson and Seymour [68] gave theorems of both the types mentioned above about these tangles: a 'tangle-tree duality' theorem and a 'tree-of-tangles' theorem. The first one, like Menger's theorem, characterizes the existence of a highly connected substructure, in that it shows that a graph contains no k-tangle if and only if the graph has branch-width at most $k - 1$.[1]

The second of these two, like the traditional theorem by Tutte, allows the

[1]branch-width is a width-measure defined via a so called *branch-decomposition*. It is closely related to tree-width in that the tree-width and the branch-width of a graph only differ by a linear factor.

decomposition of a graph into its ‘highly-connected’ pieces: there is a tree-decomposition of the graph that distinguishes all the distinguishable tangles in that they are contained in distinct bags of this decomposition. More precisely, this theorem reads as follows:

Theorem 1.1.1 ([68])**.** *Every graph has a tree-decomposition displaying its maximal tangles.*

Since neither does the existence of a k-tangle imply the existence of a k-block, nor vice versa[2], the two distinct notions of ‘highly-connected’ pieces are independent, and consequently, a unified decomposition; one distinguishing all the tangles and all the blocks simultaneously, would be a closer representation of a decomposition into ‘highly-connected pieces’. Such a decomposition was found by Diestel, Hundertmark and Lemanczyk by inventing the notion of a *profile* [27]. Profiles are, in a sense, a weaker version of a tangle in that a k-profile is again an orientation of the separations of order less than k, but with a weaker consistency condition. Diestel, Hundertmark and Lemanczyk not only showed the existence of a decomposition distinguishing all the profiles in a graph, but even were able, unlike Robertson and Seymour in Theorem 1.1.1, to do so in a canonical way: their decomposition is invariant under automorphisms of the graph at hand.

Moreover, all they used in their proof was the relation between the separations of a graph, apart from this they made no references to the graph structure at all. Consequently, their theorem is formulated in a new, abstract framework: ‘abstract separation systems’. They were able to formulate their theorem just in terms of a partially ordered set $\vec{S}$ together with an order-reversing involution * on $\vec{S}$. Here, the elements of $\vec{S}$ correspond to the oriented separations of a graph. They observed that the set of all orientated separations of a graph is a lattice, which they, in the abstract, called a *universe of separations*. So a universe of separations is a lattice $\vec{U}$ together with an order reversing involution on $\vec{U}$. Now given any separation system $\vec{S}$ contained in such a universe, a *profile* P of $\vec{S}$ is simply a way to choose exactly one orientation from $\{\vec{s}, \vec{s}^*\}$ for every separation $\vec{s} \in \vec{S}$ in a consistent way: so that if P contains a separation $\vec{s}$, and $\vec{r}$ is another separation such that $\vec{r}^* \leqslant \vec{s}$ then $\vec{r} \notin P$, and so that, whenever $\vec{r}, \vec{s}$ are contained in P, then the inverse $(\vec{r} \vee \vec{s})^*$ of their supremum, taken in U, is not contained in P. The *order* of the separations of a graph then corresponds to some symmetric submodular function on $\vec{U}$ and a k-profile, which generalizes both the k-blocks and the k-tangles mentioned above is just a profile of $\vec{S}_k$, the set of all those separations from $\vec{U}$ of order less than k. Consequently, their theorem reads as follows:

Theorem 1.1.2 (Canonical tree-of-tangles theorem for separation universes [27, Theorem 3.6])**.** *Let $U = (\vec{U}, \leqslant, ^*, \vee, \wedge, |\cdot|)$ 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 $\mathcal{P}$;

[2] although it can be shown that a k-block does give rise to a $\frac{2k}{3}$-tangle

(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.

This general abstract setup allows the application of the theory of tangles to a lot of other contexts such as matroids or arbitrary data sets.

Perhaps surprisingly, Diestel and Oum [29, 30] showed that not only a tree-of-tangles theorem can be formulated in this abstract context, but one can even prove an abstract version of a tangle-tree duality theorem, which then again could be applied to a lot of contexts other than graphs.

Consequently, the focus of tangle-theory shifted from the concrete graphs to this more abstract setup of abstract separation systems. This led to Diestel, Erde and Weißauer showing in [26] that both, a tree-of-tangles and a tangle-tree duality theorem, can in fact be formulated without the previously required submodular order function mentioned above. Instead, they distilled the properties of such an order function needed in the proof into a structural property, called structural submodularity. However, their tree-of-tangles theorem, unlike the original one for profiles from [27] does no longer build a canonical decomposition and also, since they do not require the existence of a submodular order function, they cannot guarantee that the decomposition is 'efficient' in the sense that the distinct profiles are always distinguished by separations of the lowest possible order. Their theorem reads as follows:

Theorem 1.1.3 ([26, Theorem 6])**.** *Let $\vec{S}$ be a structurally submodular separation system and $\mathcal{P}$ a set of profiles of S. Then S contains a nested set that distinguishes $\mathcal{P}$.*

1.2 The contributions of this thesis

Building on this existing theory we start with the new results presented in this thesis. Most of the work in this thesis can also be found in the papers [24, 25, 36–43], other results will be indicated explicitly.

After giving a formal overview about the notation, definitions and existing results in Chapter 2, we start in Chapter 3 with studying the following general considerations: while capturing the highly connected pieces of a graph indirectly via its tangles has a lot of advantages, it comes at the cost of no longer being able to easily describe these highly connected regions. The existence of a k-tangle in a graph does, a priori, not give us any information about that region. Consequently, in Chapter 3 we try, from different angles, to find concrete witnesses for tangles and other 'highly-connected pieces'.

We start in Section 3.1, which is based on [37], with the following: while in fact all the prototypical examples of a tangle, like for example a large grid, have the property that there exists a vertex set X (in the context of the grid, the vertex set of that grid) such that every separation in the tangle is oriented towards the side containing more vertices from X, it is an open question by Diestel [27] whether such a vertex set, which we call a decider set, exists for all tangles in graphs. Our first main theorem from Section 3.1 provides a partial solution to this:

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 $U \subseteq V$.*

While this theorem does not give us such a vertex set X as mentioned above, it at least gives us a weighted version of such a set which as such guarantees that a k-tangle in a graph at least 'points to something'. Perhaps surprisingly, this result extends to profiles of graphs and hypergraphs as well as to a variation of tangles in a graph which deals with separations corresponding to edge cuts:

Theorem 2. *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 $w(A) < w(B)$.*

However, we also show in Section 3.1 that an extension to a version of these edge-tangles for hypergraphs is not possible, by giving an example of such a tangle which does not have such a weighted decider, and thus also does not have a decider set.

After that, we deal in Section 3.2, which is based on parts of [25], with other ways to guarantee the existence of a decider, even if the considered setup is not covered by Theorem 1 or Theorem 2, because we work, for example, with the edge-tangles of hypergraphs. We obtain a sufficient condition on a tangle, which we call a high enough *resilience*, which guarantees the existence of a weighted decider for any tangle of set separations, independent of the concrete setup. More specifically, we can show the following:

Theorem 3. *Let $\vec{U}$ be the universe of all set separations of some finite set V, and let τ be an orientation of some set $S \subseteq U$ of separations. Let m be the number of maximal elements of τ. If τ is k-resilient for some $k > \frac{m}{2}$, then τ has a decider.*

The other half of Section 3.2 is devoted to finding actual decider sets, rather than just weighted versions of them. While we cannot show that every tangle, even in a graph, has a decider set, we can at least show that a restriction of the tangle to a lower order has. More precisely, we can show the following:

Theorem 4. *Let $\vec{U}$ be a universe of set separations of some finite set V equipped with the order function $|(A, B)| = |A \cap B|$ and let $k \in \mathbb{N}$. If τ' is a k-profile in $\vec{U}$ which extends to a regular $2k$-profile τ in $\vec{U}$, then τ' has a decider set $X \subseteq V$ of size $|X| \geqslant 2k$.*

Another possible way to guarantee the existence of a tangle of a separation system is via the existence of a tangle of a certain 'dual' separation system. In Section 3.3, which is based on [24], we describe a naturally arising setup of two separation systems of which we think as dual. In that setup, a tangle of one of the separation systems naturally gives an orientation of the other one. We will show that this orientation, restricted to a lower order, actually will be a tangle, i.e. we show the following:

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

In some sense, this tangle is then witnessed by the tangle of the other, 'dual' separation system.

After proving Theorem 5 which says that the 'shift' of a tangle of one separation system gives, restricted to a lower order, again a tangle, a natural question to ask is whether performing these shifting operation a second time gives a restriction of the original tangle. As it turns out, this is indeed the case, more precisely we can show the following:

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

We will also consider in Section 3.3 some variations and generalizations of this idea of obtaining one tangle from another via a 'shifting' operation.

We end Chapter 3 with a result exclusive to this thesis in Section 3.4, where we consider a different type of dense structures in graphs: 'agile sets'. Weißauer [76] proposed them as a possible generalization of the special case of 'Z-linkages' for sets Z of size 4, which were considered independently by Seymour in [70] and Thomassen in [72]. Weißauer defined a vertex set X in a graph G to be *agile* if there are, for any partition $X = X_1 \cup X_2$ of X into disjoint sets, two distinct trees T_1, T_2 in G such that X_1 is contained in T_1 and X_2 is contained in T_2. In that sense, every two partition classes of X can be connected independently. Thus, a large agile set is another object of which we might think of as highly connected. These large agile sets are not directly connected to profiles or tangles, as for example the bipartite graph $K_{2,k}$ does contain an agile set of size k, but no l-profile for $l \geqslant 3$. Weißauer asked whether, in a sense, these agile sets actually measure high connectivity in terms of these bipartite graphs by asking whether we can force the existence of a $K_{2,k}$-minor in a graph by requiring the existence of a large agile set.

In Section 3.4 we show that the general answer to this question is no, but only because of one specific additional structure: a large regular strip, the graph obtained from two long disjoint paths $P_1 = v_1 v_2 \dots v_k$ and $P_2 = w_1 w_2 \dots w_k$ by adding all the diagonal edges $v_i w_{i+1}$ and $w_i v_{i+1}$. Concretely, building on the characterization of large $K_{2,k}$-minor free graphs by Ding [32] we can show the following:

Theorem 7. *There exists a function $f : \mathbb{N} \to \mathbb{N}$, such that every graph with an agile set of size $f(k)$ either contains $K_{2,k}$ or a regular strip of length k as a minor.*

After proving this, we consider a possible generalization of agile sets. Instead of partitioning X into just two sets we might allow a partition into at most m classes, say, and consequently require that for any partition $X = X_1 \cup \dots \cup X_m$ we find disjoint trees $T_1, \dots, T_m$ so that $X_i \subseteq T_i$. This notion is what we call m-agile. As it turns out, this time using a theorem by Geelen and Joeris from [50] which generalizes the grid theorem, large m-agile sets are quantitatively characterized by the existence of either a $K_{m,N}$ or a large rectangular grid minor:

Theorem 8. *There is a function $f\colon \mathbb{N}^2 \to \mathbb{N}$ such that every graph containing an $((m-1)2m+1)$-agile set of size at least $f(m,k)$ contains $K_{m,k}$ or the $(2m-1) \times k$-grid as a minor.*

Note however, that this theorem is quantitative not only in terms of the size of the m-agile set, but also in terms of m: a $K_{m,k}$ and a $(2m-1) \times k$-grid both only contain a large m-agile set, not a large $((m-1)2m+1)$-agile set.

In Chapter 4 we then return to tangles and profiles in abstract separation systems and focus on generalizations of the existing tree-of-tangle theorems. Prior to this thesis, there were two concurring tree-of-tangles theorems, both of which can claim for themselves to be the one of greatest generality. On the one hand we have Theorem 1.1.2 by Diestel, Hundertmark and Lemanczyk from [27], on the other hand we have Theorem 1.1.3 by Diestel, Erde and Weißauer from [26].

Since a structurally submodular separation system is a weaker object than a submodular universe of separations, Theorem 1.1.3 is more widely applicable. On the other hand, Theorem 1.1.2 improves over Theorem 1.1.3 in terms of the strength of the statement: not only is the tree of tangles achieved by Theorem 1.1.2, unlike the one from Theorem 1.1.3, canonical, i.e. invariant under isomorphisms, it also distinguishes the used profiles efficiently, that is with a separation of lowest possible order. Since Theorem 1.1.3 does not require the existence of an order function, it cannot achieve this. Consequently, it would be nice to have a theorem combining the strengths of these two: it should be applicable even without an order function, but if something like an order function exists, we should be able to obtain a tree of tangles which indeed distinguishes any two profiles by a separation of lowest possible order. This is the first goal we achieve in Section 4.1, which is based on [39]. There, we are able to show the following:

Theorem 9. *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 $\mathcal{P}$.*

Not only does Theorem 1.1.3 follow from Theorem 9 by applying it to a sequence $\mathcal{S}$ consisting of just one S_1, also a non-canonical version of Theorem 1.1.2 can be obtained from Theorem 9 as a simple corollary.

But even more interesting than the result of Theorem 9 is the way we proved it. We observed that for proving tree-of-tangles theorems we actually do not even need to keep the information about the considered profiles, actually all we need is an information about the separations used to distinguish them. More precisely, we only need to consider, for every pair P, P' of profiles we want to distinguish, the set $\mathcal{A}_{P,P'}$ of separations which (efficiently) distinguish them. The information that the considered objects are actually profiles is then only used to conclude one simple property of the relation of the sets $\mathcal{A}_{P,P'}$ to one another, which we call *splinters*. The main ingredient of the proof of Theorem 9 is then our following key lemma:

Lemma 10 (Splinter Lemma)**.** *Let U be a universe of separations and let $\mathfrak{A} = (\mathcal{A}_i)_{i \leqslant n}$ be 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.*

Although the proof of Lemma 10 only takes a little more than half a page, it is the key tool to obtain various tree-of-tangles theorems. In particular, all that is required to prove the reader's favourite tree-of-tangles theorem is to check that corresponding sets $\mathcal{A}_{P,P'}$ satisfy our splinter condition, whose definition is simple and which is easy to check. We show this in Section 4.1.3 by applying Lemma 10 in a variety of contexts.

The one remaining flaw of Theorem 9 compared to the existing Theorem 1.1.2 is the fact that our splinter Lemma 10 and thus also Theorem 9 cannot achieve a canonical tree of tangles. This is due to the fact that our splinter condition is formulated as the weakest possible condition to guarantee the existence of a tree of tangles, and therefore is too weak to guarantee that the construction of such a tree can be carried out canonically. Consequently, we obtain in Section 4.1.4 a stronger, splinter-like condition which we call *splinters hierarchically* and which does allow us to obtain a canonical version of Lemma 10. Namely, we can show the following:

Lemma 11 (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 we have that $N(\varphi(\mathfrak{A})) = \varphi(N(\mathfrak{A}))$.

This lemma is again applicable in a variety of contexts, as we demonstrate in Section 4.1.5.

In Section 4.2, which, at the time of writing, is exclusive to this thesis, we are concerned with trees of tangles in one specific application: directed graphs. For the corresponding notion of a separation in a directed graph, which we call a directed separation, it is in general not possible to obtain a nested set of separations distinguishing the corresponding (directed) tangles, due to the fact that they do not allow the building of new separations out of existing ones in the same way as ordinary separations in graphs do. Consequently, they do not form a universe of separations and thus our splinter Lemma 10 cannot directly be applied to them. The best distinguishing structure for directed graphs we can hope for thus needs to be weaker, for example a *tree-labelling* as defined and constructed by Giannopoulou, Kawarabayashi, Kreutzer and Kwon in [51].

We would like to obtain a result like theirs using the ideas of our splinter Lemma 10. This cannot be done by directly applying Lemma 10, as the directed separations, as mentioned above, do not form a universe of separations. But, perhaps surprisingly, we can abstract Lemma 10 even further in order to be able to apply it to directed separations. Namely, all that is needed from the universe of separations and the separation systems in Lemma 10 is the nestedness relation on the undirected separations and the existence of certain joins and meets, which is guaranteed by the fact that the universe is a lattice. This nestedness relation is just a reflexive and symmetric relation, and the properties of a separation system required to prove Lemma 10 can entirely be described via this nestedness relation. Consequently, we can give an abstract definition of our splinter condition in terms of this nestedness relation. Using this idea we can then obtain the following, abstract version of Lemma 10:

Lemma 12 (Abstract Splinter Lemma)**.** *Let $\mathcal{A}$ be a finite set and let $\sim$ be a reflexive and symmetric relation on $\mathcal{A}$. Let $\mathfrak{A} = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$ be a family of* *subsets of $\mathcal{A}$ which splinters. Then there is a set $N = \{a_1, \ldots, a_n\} \subseteq \mathcal{A}$ such that $a_i \in \mathcal{A}_i$ for every $1 \leqslant i \leqslant n$ and such that $a \sim a'$ for any $a, a' \in N$.*

Using this abstract version of Lemma 10 we are then able to develop a theory of directed separation systems in Section 4.2.3, together with a corresponding notion of a profile, in a similar fashion as the abstract separation systems generalized the separations of an undirected graph. In the context of these directed separation systems and their profiles we can then define a weaker relation, which we call weakly $\mathcal{P}$-nested, and are able to show the existence of a tree of tangles with respect to this relation:

Theorem 13. *If $\mathcal{P}$ is a set of distinguishable profiles in a directed universe $\overleftrightarrow{U}$, then there exists a weakly $\mathcal{P}$-nested set of separations which efficiently distinguishes every two profiles from $\mathcal{P}$.*

This theorem can be applied to the directed separations of a directed graph and with this application we obtain a tree-labelling like the one constructed by Giannopoulou, Kawarabayashi, Kreutzer and Kwon in [51]. In that way, we give an alternative proof of their theorem, based on our abstract splinter Lemma 12.

After this excursion to abstract directed structures, we are, once more, concerned with the existence of canonical tree-of-tangles theorems in usual abstract separation systems. Our canonical splinter Lemma 11 does not allow us to give a canonical version of Theorem 1.1.3, due to the fact that the sets $\mathcal{A}_{P,P'}$ obtained from the profiles of a structurally submodular separation system do not necessarily splinter hierarchically. Nevertheless, we can prove a canonical version of Theorem 1.1.3, and we do so in Section 4.3, which is based on [36]. Concretely, we show the following:

Theorem 14. *Let $\vec{S}$ be a structurally submodular separation system and $\mathcal{P}$ a set of profiles of S. Then there is a nested set $N = N(\vec{S}, \mathcal{P}) \subseteq S$ which distinguishes $\mathcal{P}$. This $N(\vec{S}, \mathcal{P})$ can be chosen canonically: if $\varphi\colon \vec{S} \to \vec{S'}$ is an isomorphism of separation systems and $\mathcal{P}' := \{\varphi(P) : P \in \mathcal{P}\}$ then $\varphi(N(\vec{S}, \mathcal{P})) = N(\vec{S'}, \mathcal{P}')$.*

After having dealt with distinguishing tangles in a lot of different finite contexts, we turn our attention in Section 4.4, which is based on [42], towards infinite structures. Lemma 10 does not hold in an infinite setting due to the fact that it is proved via induction, and a transfinite version of this proof would require us to perform a limit step, which unfortunately is not possible. However, we can obtain a version of Lemma 11 in the infinite setting. We this time formulate this Lemma directly in the more abstract from of Lemma 12, as we will need to apply it in this more abstract form. Moreover, our corresponding infinite version of 'splinters hierarchically', which we call 'splinters thinly' is slightly stronger as this definition additionally requires, roughly speaking, that the separation at hand only crosses with finitely many others of the same order. Apart from that, we can obtain exactly the same theorem as Lemma 11 in an infinite setting:

Lemma 15. *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 $N((\,\varphi(\mathcal{A}_i) \,:\, i \in I\,)) = \varphi(N((\,\mathcal{A}_i \,:\, i \in I\,)))$.*

Unlike for finite graphs, Lemma 15 cannot directly be applied in the context of profiles (or tangles) of an infinite (not locally finite) graph, as the considered separations there would not need to splinter thinly. This, however, is in line with the existing theory: there does not in fact exist a theorem stating that all tangles of an infinite graph can canonically be distinguished by a nested set of separations, not even if one special type of 'silly' tangles, called ultrafilter tangles, is excluded. In fact, such a nested set can only be obtained in a non-canonical way, as shown by Carmesin [9]. However, as shown by Carmesin, Hamann, Miraftab [14], there does exist a canonical structure to distinguish the tangles of an infinite graph, which they call a tree of tree-decompositions. Both these results can be deduced from Lemma 15, via an additional twist. We obtain both of these results via an intermediate step which is interesting in its one right: while it is not possible to find a canonical nested set of separations which distinguishes all the tangles of an infinite graph, it is possible to define a notion of nestedness for the *separators* of such separations for which this is possible. Namely, we are able to show the following:

Theorem 16. *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 $\mathcal{P}$.*

This nested set of separators can then be, on the one hand, used to construct a non-canonical nested set of separations like the one from [9], and on the other hand to obtain a tree of tree-decompositions as defined by Carmesin, Hamann, Miraftab.

While this application of Lemma 15 to classical separations of infinite graphs is a bit tricky, there does in fact exist a prototypical application of Lemma 15: it can be naturally applied to the *edge-blocks* of a graph, where a k-edge-block shall be a $\subseteq$-maximal subgraph which cannot be separated by a set of less than k edges. Consequently, we show in Section 4.5, which is based on [43], the following:

Theorem 17. *Every connected graph G has a nested set of bonds that efficiently distinguishes all the edge-blocks of G.*

This theorem can also be deduced from the theory of edge-cuts developed by Dicks and Dunwoody [17], as they showed that every graph contains a canonical nested set of bonds which generates all other bonds. Here, a bond can be generated from a set of bonds if the separation corresponding to that bond can be obtained from the separations corresponding to the set of bonds via finite combination of join and meet operations. As it turns out, not only does such a set of bonds which generates all bonds need to distinguish all the edge-blocks efficiently, also the other direction holds: every nested set of bonds which efficiently distinguishes all the edge blocks will also generate all bonds, i.e. we can show the following result:

Theorem 18. *Let G be any connected graph and let M be any nested set of bonds of G. Then the following assertions are equivalent:*

1. *M efficiently distinguishes all the edge-blocks of G;*
2. *For every $k \in \mathbb{N}$, the $\leqslant k$-sized bonds in M generate all the k-sized cuts of G.*

Thus, Theorem 17 and Theorem 18 together give an alternative, combinatorial proof based on Lemma 15 of the important result by Dicks and Dunwoody.

The last Section 4.6 of Chapter 4, which is based on [41], is devoted not to the development of even more general tree-of-tangles theorems, but to obtaining such a theorem via a completely different proof method. The theory of tangles consists, as already stated in the introduction, of two major types of theorems: on the one hand there are tree-of-tangles theorems which seek out to distinguish all the given tangles of our structure, on the other hand there are tangle-tree duality theorems which are concerned with structures witnessing the absence of tangles. While these two types of theorems are fundamentally different, we show in Section 4.6 that actually the most general form of a tangle-tree duality theorem can be used to prove tree-of-tangles theorems, for example a version of Theorem 1.1.3. Consequently, the two archetypical results from tangle theory are not as fundamentally different, as one might think.

We prove a tree-of-tangles theorem via a tangle-tree duality theorem by considering a hand crafted structure as a 'tangle', which cannot exist in the given structure. Consequently, the tangle-tree duality theorem will then give us a tree-like structure witnessing the absence of this type of tangle. This structure can then be used to obtain the desired tree-of-tangles theorem. Proving a tree-of-tangle theorem in this way actually has the benefit of allowing us, in some cases, to give some bound on the degree of the nodes in such a tree, something which is not so easy to do with the conventional methods. For example, we can show the following:

Theorem 19. *Let $\vec{U}$ be a submodular universe and let $\mathcal{P}$ be the set of regular profiles of $\vec{S}_k$. Then there exists a 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 $\max\{\delta_e(\mathcal{P}),3\}$.*

The last Chapter 5 of this thesis deals with abstract separation systems as an object of its own right. While a lot of results in this thesis are formulated in the context of these separation systems, not so much is known about their general properties. For example, one question we answer in Section 5.1, which is based on [38], is the following: prior to this thesis, it was not known whether there actually exists a submodular separation system which does not at all come from a submodular order function. If in fact all such systems would come from some submodular order function, then all the results about structurally submodular separation systems would actually be just corollaries from the corresponding result about separation systems with an order function. But, as it turns out, there indeed are such examples and we construct one in Section 5.1.3:

Theorem 20. *There exists a separation system $\vec{S}$ which is submodular in a universe $\vec{U}$ of set bipartitions whose submodularity in $\vec{U}$ is not induced by a submodular order function on $\vec{U}$.*

Another question about these abstract separation systems deals with their interior structure. It has been characterized [3] which abstract separation systems do arise as the special type of bipartitions of a finite set. What we are able to show in Section 5.1.5 is that in fact every structurally submodular separation system, in a distributive universe, can be built from these systems of bipartitions. More specifically, we can show the following:

Theorem 21. *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.*

Finally, in Section 5.2, which is based on [40], we focus again on the difference between separation systems arising from an order function and the structurally submodular ones. It can be shown that the former ones always have the following property: if S_k is the set of all separations of order less than k, then we find a separation in S_k so that after deleting this separation we are left with a separation system that is still structurally submodular. If an equivalent property would hold for every structurally submodular separation system S, then each such system could be obtained via a sequence $\emptyset = S_0 \subseteq S_1 \subseteq \cdots \subseteq S_n = S$ of structurally submodular separation systems, where any two consecutive systems differ in just one separation. We could then think of the 'order' of a separation as the smallest k for which that separation appears in S_k and would hope that some of the proofs using a submodular order function could still be carried out for this definition of order. Unfortunately, at least for separation systems in a non-distributive universe, such a sequence does not need to exist, as we show in Section 5.2.5. For separation systems in distributive universes on the other hand, the problem turns out to be equivalent to the following easy to state question about subsets of a finite set:

Problem 22 (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?

While we do not know the general answer to Problem 22, we still can provide a partial positive answer to a variation of that problem.

Instead of measuring the submodularity of a separation system 'externally', inside a surrounding universe of separations, we may as well measure this submodularity 'internally' inside the separation system itself. That is, we may say that a separation system $\vec{S}$ is *submodular* if for any two separations $\vec{s}$ and $\vec{t}$ in that system there exists a common join or a common meet inside that separation system. A similar property may also be defined just for arbitrary posets: let us call a poset P *woven* if any two elements of that poset have a join or a meet, calculated inside the poset. Now, an unravelling of such a woven poset P is a sequence $\emptyset = P_0 \subseteq \cdots \subseteq P_n = P$ as above: with the property that every P_i is woven and any two consecutive P_i differ by only one element. What we can show in Section 5.2.6 is the following:

Theorem 23. *Every woven poset can be unravelled.*

An analogue statement also holds for the 'internally' submodular separation systems mentioned above.

Chapter 2

Preliminaries

We use the basic graph-theoretic notation from [20]. Moreover, this thesis is build on the theory of (abstract) separation systems introduced for example in [21, 27]. In what follows we will recap the definitions from that context that we need and, in some cases, generalize them slightly to feed our needs.

2.1 Lattice Theory

The definitions in this section are based on [16]. A *partially ordered set* or *poset* is a set P together with a binary relation $\leqslant$ on P which is reflexive, antisymmetric, and transitive. Given a poset $(P, \leqslant)$ and $a, b \in P$ we write $a < b$ to mean that both $a \leqslant b$ and $a \neq b$ hold.

A *lattice* is a non-empty partially ordered set $(L, \leqslant)$, in which any two points a and b have a join and a meet. Here, a *join* (or *supremum*) of a and b is an element $c \in L$ such that $a \leqslant c$ and $b \leqslant c$ with the additional property that $c \leqslant d$ whenever $a \leqslant d$ and $b \leqslant d$. So there is a *smallest upper bound* for a and b. This element is then denoted as $a \vee b$. Analogously, a *meet* (or *infimum*) of a and b is an element $c \in L$ such that $c \leqslant a$ and $c \leqslant b$ with the additional property that $d \leqslant c$ whenever $d \leqslant a$ and $d \leqslant b$. So there is a *largest lower bound* for a and b. This element is then denoted as $a \wedge b$.

Every finite lattice $(L, \leqslant)$ has a *top element* $\top$ satisfying $a \leqslant \top$ for all $a \in L$ and a *bottom element* $\bot$ satisfying $\bot \leqslant a$ for all $a \in L$.

A lattice $(L, \leqslant)$ is said to be *distributive* if the distributive laws hold, i.e. $a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c)$ and $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$ for all $a, b, c \in L$.

An isomorphism between two lattices $(L, \leqslant)$ and $(L', \leqslant')$ is a bijection $\varphi : L \to L'$ respecting the partial orders, that is for $a, b \in L$ we have that $a \leqslant b$ if and only $\varphi(a) \leqslant \varphi(b)$. Note that such an isomorphism needs to respect joins and meets, i.e. $\varphi(a \vee b) = \varphi(a) \vee \varphi(b)$ and $\varphi(a \wedge b) = \varphi(a) \wedge \varphi(b)$ for all $a, b \in L$.

A typical example of a lattice is the so called *subset lattice* of a set V: the set of all subsets of V ordered by inclusion (so $A \leqslant B$ if and only is $A \subseteq B$) is a distributive lattice, where $A \vee B = A \cup B$ and $A \wedge B = A \cap B$ for all $A, B \subseteq V$.

2.2 Abstract Separation Systems

In this section, we follow the notation from [21]. A *separation system* $(\vec{S}, \leqslant, {}^*)$ consists of a partially ordered set $(\vec{S}, \leqslant)$ together with an *order-reversing involution* *. Here, an *involution* is a map ${}^* : \vec{S} \to \vec{S}$ such that $(\vec{s}^*)^* = \vec{s}$ and * being *order-reversing* means that $\vec{s} \leqslant \vec{t}$ if and only if $\vec{t}^* \leqslant \vec{s}^*$ for all $\vec{s}, \vec{t} \in \vec{S}$. In most parts of this thesis, we will assume that the considered separation systems are finite, we will explicitly state where we do not require this. The definitions in Chapter 2 however are formulated such that they also hold for infinite separation systems.

The elements of $\vec{S}$ are called *oriented separations.* Given a separation $\vec{s}$ inside a separation system $\vec{S}$, we also write $\overleftarrow{s}$ as shorthand for the *inverse* $\vec{s}^*$ of $\vec{s}$ under *. For the set $\{\vec{s}, \overleftarrow{s}\}$ consisting of an oriented separation $\vec{s}$ together with its inverse, we use the shorthand notation s and call s an *unoriented separation* or the *underlying unoriented separation of* $\vec{s}$ (or $\overleftarrow{s}$). Moreover, we write S to denote the set of all the unoriented separations belonging to separations in $\vec{S}$. The other way around, given an unoriented separation s the *orientations* of s are the oriented separations contained in s and denoted by $\vec{s}$ and $\overleftarrow{s}$. Given a set T of unoriented separations we mean by $\vec{T}$ the set of all oriented separations contained in the elements of T.

If there is no risk of confusion we may use the term *separation* to mean either an oriented or an unoriented separation, and we will use terms and properties only defined for one of the two types also for the other one if the meaning is clear. Moreover, we may also talk about the *separation system* S to mean a separation system $(\vec{S}, \leqslant, {}^*)$ of which S is the set of unoriented separations.

An *isomorphism* between separation systems $(\vec{S}, \leqslant, {}^*)$ and $(\vec{S'}, \leqslant', {}^\circ)$ is a bijection φ between $\vec{S}$ and $\vec{S'}$ respecting the partial order and the involution, so $\vec{s}^* = \varphi(\vec{s})^\circ$ and $\vec{s} \leqslant \vec{t}$ if and only if $\varphi(\vec{s}) \leqslant' \varphi(\vec{t})$ for all $\vec{s}, \vec{t} \in \vec{S}$.

Two unoriented separations s and t from a separation system S are *nested* if there are orientations $\vec{s}$ and $\vec{t}$ of s and t such that $\vec{s} \leqslant \vec{t}$. Note that this is a symmetric property, as $\vec{s} \leqslant \vec{t}$ implies that $\overleftarrow{t} \leqslant \overleftarrow{s}$. If s and t are not nested, they *cross*, or are *crossing* separations. We will use the terms of *nested* and *crossing* separations also for oriented separations to mean that the unoriented separations corresponding to them cross.

We say that two oriented separations $\vec{s}$ and $\vec{t}$ *point towards each other* if $\vec{s} \leqslant \overleftarrow{t}$. Note that this implies that also $\vec{t} \leqslant \overleftarrow{s}$. The separation $\vec{s}$ and $\vec{t}$ *point away from each other* if $\overleftarrow{t} \leqslant \vec{s}$, and thus also $\overleftarrow{s} \leqslant \vec{t}$. So, if two oriented separations $\vec{s}$ and $\vec{t}$ are nested, they either are comparable (so $\vec{s} \leqslant \vec{t}$ or $\vec{t} \leqslant \vec{s}$), or they point towards each other, or they point away from each other.

A set N of unoriented separations is *nested* if the elements of N are pairwise nested.

A separation $\vec{s}$ in $\vec{S}$ is called *small* if $\vec{s} \leqslant \overleftarrow{s}$. In that case we say that $\overleftarrow{s}$ is *co-small.* A set $\vec{N}$ of oriented separations is *regular* if $\vec{N}$ does not contain any co-small separations. Consequently, a set N of unoriented separations is *regular* if $\vec{N}$ is regular, i.e. if no separation in N has a small orientation.

A separation $\vec{s}$ in $\vec{S}$ is called *trivial in* $\vec{S}$ if there is a separation t in S such that $\vec{s} < \vec{t}$ and $\vec{s} < \overleftarrow{t}$. As above, in that case we say that $\overleftarrow{s}$ is *co-trivial in* $\vec{S}$ and that s is *trivial in* S. Note that every trivial separation is small, but the opposite is not true in general. In particular, whether a separation is trivial or not depends on the surrounding separation system.

Given a separation system S, a nested set N of unoriented separations from S is called a *tree set* if N does not contain any separation which is trivial in S.

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

A *star* of separations is a set σ of oriented separations which are not degenerate and all point towards each other. That is, any two distinct separations $\vec{s}, \vec{r} \in \sigma$ satisfy that $\vec{s} \leqslant \overleftarrow{r}$. In particular, any two separations in a star are nested.

2.3 Universes of separations

Like for the previous section, the terms in this section are based on [21].

If a separation system $(\vec{U}, \leqslant, ^*)$ happens to be a lattice we say that U or $(\vec{U}, \leqslant, ^*)$ is a *universe of separations*. In universes, DeMorgan's law holds, i.e. $(\vec{s} \vee \vec{t})^* = \overleftarrow{s} \wedge \overleftarrow{t}$.

A *homomorphism between universes* $(\vec{U}, \leqslant, ^*)$ and $(\vec{U'}, \leqslant', ^\circ)$ is any map $\varphi : \vec{U} \to \vec{U'}$ which respects the involutions and the joins and meets of $\vec{U}$ and $\vec{U'}$, so $\varphi(\overleftarrow{s}) = \varphi(\vec{s})^\circ$, $\varphi(\vec{s} \vee \vec{t}) = \varphi(\vec{s}) \vee \varphi(\vec{t})$, and $\varphi(\vec{s} \wedge \vec{t}) = \varphi(\vec{s}) \wedge \varphi(\vec{t})$ for all $\vec{s}, \vec{t} \in \vec{U}$. Note that each such homomorphism respects the partial order on $\vec{U}$, as given $\vec{s} \leqslant \vec{t}$ in $\vec{U}$ we have that $\varphi(\vec{t}) = \varphi(\vec{s} \vee \vec{t}) = \varphi(\vec{s}) \vee \varphi(\vec{t})$ and thus $\varphi(\vec{s}) \leqslant \varphi(\vec{t})$.

Observe that, if U and U' are universes and $\varphi : \vec{U} \to \vec{U'}$ is an isomorphism between the underlying separation systems $(\vec{U}, \leqslant, ^*)$ and $(\vec{U'}, \leqslant', ^\circ)$, then φ is also a homomorphism between U and U', i.e. φ respects the join and meet operations. Consequently, we say that such a map is an *isomorphism of universes*. Note that if φ is an isomorphism of universes, then φ also is an isomorphism between the underlying lattices. Like for separation systems, we say that two universes are *isomorphic* if there is such an isomorphism.

Given two unoriented separations s and t inside a universe any separation of the form $\vec{s} \vee \vec{t}$ or $\vec{s} \wedge \vec{t}$ is called a *corner* or *corner separation* of s and t. Moreover, we will also call any underlying unoriented separation of such a corner separation a corner (or corner separation) of s and t. Consequently, the set of unoriented corner separations consists of four separations (although it is possible that some of these unoriented separations may coincide).

An important basic property of these corner separations is collected in the so called *fish lemma*:

Lemma 2.3.1 ([21, Lemma 3.2])**.** *Let $r, s \in U$ be two crossing separations. Every separation t that is nested with both r and s is also nested with all four corner separations of r and s.*

2.4 Submodularity

The notation in this section is also based on [21]. Usually (except for some part of Chapter 5) we are interested in separation systems which are contained in some surrounding universe of separations and fulfil some kind of submodularity inside this universe.

One possible variant of this submodularity is the existence of a *submodular order function* on $\vec{U}$. Here, an *order function* is a function from the set $\vec{U}$ into the natural numbers[1], i.e. $|\cdot| : \vec{U} \to \mathbb{N}$ which is *symmetric* in that $|\vec{s}| = |\overleftarrow{s}|$ for all

[1] In this thesis, the set $\mathbb{N}$ of natural numbers includes 0.

$\vec{s} \in \vec{U}$. We then also write $|s|$ to mean $|\vec{s}| = |\overleftarrow{s}|$. That an order function $|\cdot|$ is *submodular* means that $|\vec{s}| + |\vec{t}| \geqslant |\vec{s} \vee \vec{t}| + |\vec{s} \wedge \vec{t}|$ for all $\vec{s}, \vec{t} \in \vec{U}$.

We say that a universe U is *submodular* or a *submodular universe* if it comes with a submodular order function. Given such a submodular universe, we usually consider the separation systems $\vec{S}_k := \{\vec{s} \in \vec{U} \ : \ |\vec{s}| < k\}$ of all separations of order less than k, for k some integer or $k = \aleph_0$. Note that $\vec{S}_{\aleph_0} = \vec{U}$.

In some parts of the literature such a submodular order function is also considered to be a function $f : \vec{U} \to \mathbb{R}$ into the reals instead of the natural numbers. This however only is the case in some parts of the literature which deal with finite universes, and in finite universes, due to the fact that the reals can be considered as a vector space over the rationals and moreover the rationals are dense in the reals, every such function into the reals can be replaced by a function into the natural numbers, without changing the family of the sets S_k, as shown in the following lemma:

Lemma 2.4.1. *Let U be a finite universe and let $f : \vec{U} \to \mathbb{R}$ be a function that is symmetric and submodular. Then there is a submodular order function $|\cdot| : \vec{U} \to \mathbb{N}$ such that there is, for every real number b, a natural number k such that $\{\vec{s} \in \vec{U} \ : \ f(\vec{s}) < b\} = \vec{S}_k$ and conversely, for every natural number k there is a real number b such that $\{\vec{s} \in \vec{U} \ : \ f(\vec{s}) < b\} = \vec{S}_k$.*

Proof. We will first show that we can find a rational valued function with which we can replace f. For this we note that, since U is finite, there are only finitely many distinct values $a_1 < a_2 < \cdots < a_n$ the function f takes. Moreover, there exists an $\epsilon > 0$ such that submodularity is always satisfied with either equality, or with a difference of size at least ϵ and that moreover the difference between any two a_i is at least ϵ. For example, we can choose ϵ so that for any $a, b, c, d \in \{a_1, \dots, a_n\}$ we either have $a + b = c + d$ or $|a + b - c - d| > \epsilon$.

We now choose a collection $\{b_1, \dots, b_k\}$ of real numbers so that $\{b_1, \dots, b_k\}$ is $\subseteq$-minimal with the property that there are, for every a_i, rational numbers $q_{1,i}, \dots, q_{k,i}$ such that $a_i = \sum_{j=1}^{k} q_{j,i} b_j$. Since $\{a_1, \dots, a_n\}$ is a candidate for this collection, such a collection exists.

Now we choose, for $1 \leqslant j \leqslant k$, the number m_j to be the maximal $|q_{j,i}|$, where this maximum is taken over all i from 1 to n. Moreover, let ϵ_j be smaller than $\frac{\epsilon}{4km_j}$. Since $\mathbb{Q}$ is dense in $\mathbb{R}$, we find rational numbers r_j so that $|b_j - r_j| < \epsilon_j$ for every $1 \leqslant j \leqslant k$. We now construct a function $f_q : \vec{U} \to \mathbb{Q}$ by declaring that $f_q(\vec{s}) = \sum_{j=1}^{k} q_{j,i} r_j$ whenever $f(\vec{s}) = a_i$.

Clearly, f_q is symmetric, we claim that f_q is again submodular and that there is, for every real number b, a rational number q such that

$$\{\vec{s} \in \vec{U} \ : \ f(\vec{s}) < b\} = \{\vec{s} \in \vec{U} \ : \ f(\vec{s}) < q\}.$$

To see that f_q is submodular let $\vec{s}, \vec{t} \in \vec{U}$. If $f(\vec{s}) + f(\vec{t}) \neq f(\vec{s} \vee \vec{t}) + f(\vec{s} \wedge \vec{t})$ we know, by the choice of ϵ, that $f(\vec{s}) + f(\vec{t}) \geqslant f(\vec{s} \vee \vec{t}) + f(\vec{s} \wedge \vec{t}) + \epsilon$. However, by our choice of the r_j we know that $|f_q(\vec{u}) - f(\vec{u})| \leqslant \frac{\epsilon}{4}$ for any $\vec{u} \in \vec{U}$, thus indeed $f_q(\vec{s}) + f_q(\vec{t}) \geqslant f_q(\vec{s} \vee \vec{t}) + f_q(\vec{s} \wedge \vec{t})$.

If on the other hand $f(\vec{s}) + f(\vec{t}) = f(\vec{s} \vee \vec{t}) + f(\vec{s} \wedge \vec{t})$, say because of $f(\vec{s}) = a_w, f(\vec{t}) = a_x, f(\vec{s} \vee \vec{t}) = a_y, f(\vec{s} \wedge \vec{t}) = a_z$ and $a_w + a_x = a_y + a_z$, we claim that $q_{j,w} + q_{j,x} = q_{j,y} + q_{j,z}$ for every $1 \leqslant j \leqslant k$. Indeed, we know that

$$\sum_{j=1}^{k} q_{j,w} b_j + \sum_{j=1}^{k} q_{j,x} b_j = \sum_{j=1}^{k} q_{j,y} b_j + \sum_{j=1}^{k} q_{j,z} b_j$$

and if $q_{m,w} + q_{m,x} \neq q_{m,y} + q_{m,z}$, say, then

$$b_m = \frac{\sum_{j=1,j\neq m}^{k} q_{j,w} b_j + \sum_{j=1,j\neq m}^{k} q_{j,x} b_j - \sum_{j=1,j\neq m}^{k} q_{j,y} b_j - \sum_{j=1,j\neq m}^{k} q_{j,z} b_j}{q_{m,y} + q_{m,z} - q_{m,w} - q_{m,x}}$$

which contradicts the choice of $\{b_1, \ldots b_l\}$, as this means that b_m is a rational linear combination out of the other b_j, hence $\{b_1, \ldots, b_{m-1}, b_{m+1}, \ldots, b_k\}$ contradicts the $\subseteq$-minimality of $\{b_1, \ldots b_l\}$.

Thus, indeed $q_{j,w} + q_{j,x} = q_{j,y} + q_{j,z}$ for every $1 \leqslant j \leqslant l$ and hence we also have, by the construction of f_q, that $f_q(\vec{s}) + f_q(\vec{t}) = f_q(\vec{s} \vee \vec{t}) + f_q(\vec{s} \wedge \vec{t})$.

Since the value of f_q only depends on the value of f, we clearly also have that $f_q(\vec{s}) = f_q(\vec{t})$ whenever $f(\vec{s}) = f(\vec{t})$. Moreover, if $f(\vec{s}) < f(\vec{t})$, say, then clearly our choice of ϵ implies that $f(\vec{s}) + \epsilon \leqslant f(\vec{t})$, and thus, since $|f(\vec{s}) - f_q(\vec{s})| \leqslant \frac{\epsilon}{4}$, we have that also $f_q(\vec{s}) < f_q(\vec{t})$. Thus, it is easy to find for every real number b a rational number r so that $\{\vec{s} \in \vec{U} : f(\vec{s}) < b\} = \{\vec{s} \in \vec{U} : f_q(\vec{s}) < r\}$ and conversely, for every rational number r to find a real number b so that $\{\vec{s} \in \vec{U} : f(\vec{s}) < b\} = \{\vec{s} \in \vec{U} : f_q(\vec{s}) < r\}$.

Now to obtain an actual submodular order function out of f_q we observe that, again since $\vec{U}$ is finite, there is a natural number M such that Mf_q is a function with values in $\mathbb{Z}$. Now let N be the minimal value of Mf_q then $|\vec{s}| = M \cdot f_q(\vec{s}) - N$ is a function with values in $\mathbb{N}$, and it is easy to show that this is the required order function. □

In the infinite setting however, a statement like Lemma 2.4.1 is no longer possible. Since there the theory only works for an order function which indeed takes it values in the natural numbers, order functions in this thesis are always required to be $\mathbb{N}$-valued. However, all statements about finite universe with an order function can, using Lemma 2.4.1, be easily translated to also work if instead a real-valued symmetric submodular function is given.

While all the theory of tangles was originally performed in these submodular universes, nowadays one usually can work with a weaker property. Namely, with the notion of (structurally) submodular separation systems inside a universe, developed by Diestel, Erde and Weißauer in [26]. Given a universe U, we say that a separation system $\vec{S} \subseteq \vec{U}$ is *(structurally) submodular in* $\vec{U}$ if for any two $\vec{s}, \vec{t} \in \vec{S}$ we have that $\vec{s} \vee \vec{t} \in \vec{S}$ or $\vec{s} \wedge \vec{t} \in \vec{S}$, where $\vec{s} \vee \vec{t}$ and $\vec{s} \wedge \vec{t}$ are the joins and meets taken in the surrounding universe $\vec{U}$. Note that this is not the same as requiring that there is a join or a meet of $\vec{s}$ and $\vec{t}$ in the poset $\vec{S}$, this is another possible notion of submodularity which is considered in parts of Chapter 5.

The notion of structural submodularity of S in U is weaker than requiring that U is submodular in that, given a submodular universe U, every S_k needs to be structurally submodular in U, due to the fact that the order function on U is submodular.

2.5 Orientations

The definitions in this section are from [27]. Given a separation system S, a subset $O \subseteq \vec{S}$ is *antisymmetric* if $|O \cap \{\vec{s}, \overleftarrow{s}\}| \leqslant 1$ for every separation $\vec{s} \in \vec{S}$. In that case we say, given some $s \in S$, that O *orients* s *(as* $\vec{s}$*)* if $\vec{s} \in O$.

If O contains exactly one orientation of every separation in S, that is $|O \cap \{\vec{s}, \overleftarrow{s}\}| = 1$ for every $\vec{s} \in \vec{S}$, i.e. O orients every separation from S, then O is said to be an *orientation of* S.

Such an orientation O is *consistent* if O does not contain any two distinct separations $\vec{r}$ and $\vec{s}$ such that $\overleftarrow{r} \leqslant \vec{s}$. In particular, a consistent orientation of S cannot contain any separation which is co-trivial in $\vec{S}$, as if $\vec{s} \in \vec{S}$ is co-trivial, witnessed by the separation $t \neq s$, then $\vec{t} \leqslant \vec{s}$ and $\overleftarrow{t} \leqslant \vec{s}$ and thus a consistent orientation containing $\vec{s}$ can neither contain $\vec{t}$ nor $\overleftarrow{t}$, i.e. would not orient t.

A separation s (or its orientations) *distinguishes* two orientations O and O' (of potentially different separation systems) if s is not degenerate and one of the two orientations contains $\vec{s}$ and the other contains $\overleftarrow{s}$.

Two orientations of (potentially different) separation systems are *distinguishable* if there is a separation distinguishing them.

This notation naturally extend to sets of orientations, a set $\mathcal{O}$ of orientations is *distinguishable* if any two orientations in that set are distinguishable, and a set N of separations *distinguishes* $\mathcal{O}$ if any two orientations in $\mathcal{O}$ are distinguished by a separation in N. Note that no separation distinguishing two consistent orientations of the same separation system S can be trivial or co-trivial in S.

If we are given a submodular universe U and two orientations O and O' of S_k and $S_{k'}$, respectively, where $k, k' \in \mathbb{N} \cup \{\aleph_0\}$, then a separation $s \in U$ distinguishes O and O' *efficiently* if s has minimal possible order among all separations distinguishing O and O'.

One special type of consistent orientations are the profiles. Given a separation system S that is contained in some universe U, we say that a consistent orientation P of S is a *profile of* S if P satisfies the *profile property* (P):

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

Note that here the meet is again taken in U. Further recall that a profile is *regular* if it does not contain any co-small separations.

If we consider a submodular universe U, so a universe with a submodular order function, a profile of S_k, for some $k \in \mathbb{N} \cup \{\aleph_0\}$, is said to be a *$k$-profile in* U. If the k is not important we might omit it and say that P is a *profile in* U to mean that there is some k such that P is a profile of S_k. In that case we say that k is the *order of* P, note that if there are k and k' such that $\vec{S}_k = \vec{S}_{k'}$, any k-profile is also a k'-profile and thus a profile may have multiple orders.

A k-profile P is said to be *robust* if

$$\forall \vec{s} \in P, \vec{t} \in \vec{U}: \text{ if } |\vec{s} \vee \vec{t}| < |\vec{s}| \text{ and } |\vec{s} \vee \overleftarrow{t}| < |\vec{s}|, \text{ then either } \vec{s} \vee \vec{t} \in P \text{ or } \vec{s} \vee \overleftarrow{t} \in P.$$

Note that this robustness property is usually only needed for separations $\vec{t}$ that lie in a slightly smaller subset than the whole of $\vec{U}$. This is why Diestel, Hundertmark and Lemanczyk [27] defined the slightly weaker, but also much more technical notion of profiles being *n-robustly distinguishable* and declared that a *robust set of profiles* shall be a set of profiles that is n-robustly distinguishable for a large enough n. Since the notion of robustness only plays a minor role in this thesis, we are not going to define this slightly more general notion formally. Instead, the statements in this thesis will be formulated for a *set of robust profiles*. Every distinguishable set of robust profiles is a robust set of profiles. Moreover, the reader familiar with the formal definition of n-robustly distinguishable profiles may replace every occurrence of 'set of robust profiles' with 'robust set of profiles'.

2.6 Tree sets and S-trees

The definitions in this section are based on [22, 52].

Given a graph-theoretic tree T, finite or infinite, let us denote the set of orientations of its edges as $\vec{E}(T)$, i.e. $\vec{E}(T)$ contains for every edge $e = vw \in E(T)$ its two orientations (v, w) and (w, v). On this set of these orientations we have a natural partial order by defining, for $\vec{e} = (v_1, v_2)$ and $\vec{f} = (w_1, w_2)$, that $\vec{e} \leqslant \vec{f}$ if and only if the unique path in T from v_1 to w_2 contains both v_2 and w_1. Together with the involution * on $\vec{E}(T)$ which reverses the edges, i.e. $(v, w)^* = (w, v)$, this partial order turns $\vec{E}(T)$ into a separation system. In fact, since any two directed edges are nested and no directed edge is small, this separation system is a regular tree set which we call the *edge tree set of* T.

Given any separation system $\vec{S}$, an *S-tree* (T, α) consists of a (finite or infinite) tree T together with a map α from $\vec{E}(T)$ to $\vec{S}$ that commutes with * in that $\alpha(\overleftarrow{e}) = \alpha(\vec{e})^*$.

Such an S-tree (T, α) is called *order-respecting* if for any $\vec{e} \leqslant \vec{f} \in \vec{E}(T)$ we have that $\alpha(\vec{e}) \leqslant \alpha(\vec{f})$, i.e. α preserves the partial order. Given a node $t \in V(T)$ the set $\{\alpha(s, t) : s \in N(t)\}$ of all the images under α of the ingoing edges to t is denoted as $\alpha(t)$.

These S-trees are used to relate nested sets of separations to tree-like structures.

Given some (finite or infinite) separation system $\vec{S}$ and an ordinal number α (we shall only use $\alpha \in \mathbb{N}$ or $\alpha = \omega$ or $\alpha = \omega + 1$), a *chain of order type* α or an *α-chain*, or a *chain of length* α, is a collection $\{\vec{s}_i : 0 \leqslant i < \alpha\}$ of oriented separations such that $\vec{s}_i < \vec{s}_j$ whenever $i < j$. Such a chain is *contained* in a set N of unoriented separation if all the separations in that chain are orientations of separations from N.

Now Kneip and Gollin [52] used these chains to characterize which tree sets arise as the edge tree sets of trees:

Theorem 2.6.1 ([52], Theorem 3.9)**.** *A regular tree set is isomorphic to the edge tree set of a suitable tree if and only if it does not contain a chain of order type* $\omega + 1$.

This result also immediately gives the existence of a suitable order-respecting S-tree. For finite separation systems, this was shown in [22]:

Lemma 2.6.2 (see also [22])**.** *Let N be a regular tree which does not contain an $\omega + 1$-chain. Then there exists an order-respecting S-tree (T, α) such that the image of α equals $\vec{N}$.*

Proof. By Theorem 2.6.1, the set N is isomorphic to the edge tree set of a tree T. This implies that the corresponding isomorphism α is order-respecting, i.e. (T, α) is an order-respecting S-tree, with image $\vec{N}$. □

Having established this connection between tree sets and order-respecting S-trees, we would like to be able to explicitly characterize how the S-tree corresponding to a tree set can be constructed. As it turns out, in the context were we need this, this is indeed possible: namely we will usually use tree sets to distinguish orientations and, if a tree set is chosen $\subseteq$-minimal with keeping this property, we can explicitly construct the corresponding S-tree. For finite separation systems in graphs, something like this was shown in [12], and our proof is somewhat similar. For showing this we first need to observe how $\subseteq$-minimal nested sets of separations behave:

Lemma 2.6.3. *Let $\mathcal{O}$ be a set of consistent orientations of some separation system S (finite or infinite) and let $N \subseteq S$ be a $\subseteq$-minimal nested set with the property that N distinguishes any two orientations in $\mathcal{O}$ from each other. Then there is, for every separation s in N, a unique pair O, O' of orientations from $\mathcal{O}$ such that s is the only separation in N which distinguishes O and O'.*

Proof. If for some $s \in N$ there would be no such pair, then $N \setminus \{s\}$ would still be a nested set which distinguishes $\mathcal{O}$, contradicting the $\subseteq$-minimality. Now let $s \in N$ and $O_1, O_1', O_2, O_2' \in \mathcal{O}$, such that $\vec{s} \in O_1, O_2$ and $\overleftarrow{s} \in O_1', O_2'$ and let us suppose that $O_1 \neq O_2$. Since N distinguishes all orientations in $\mathcal{O}$ from another, N needs to contain some separation $t \neq s$ which distinguishes O_1 and O_2 from another, say because of $\vec{t} \in O_1, \overleftarrow{t} \in O_2$. Since s and t are nested we either have $\vec{s} \leqslant \vec{t}$ or $\vec{s} \leqslant \overleftarrow{t}$. In the first case t distinguishes O_1 and O_1', whereas in the second case t distinguishes O_2 and O_2', in particular s is either not the only separation from N distinguishing O_1 and O_1', or not the only separation from N distinguishing O_2 and O_2'. □

Moreover, we can show that for finite separation systems the S-tree corresponding to such a $\subseteq$-minimal regular tree set is essentially unique:

Lemma 2.6.4. *Let $\mathcal{O}$ be a set of consistent orientations of some separation system S, let $N \subseteq S$ be a $\subseteq$-minimal regular tree set such that N distinguishes all the orientations from $\mathcal{O}$ from one another. Given any order-respecting S-tree (T, α) with the property that the image of α equals $\vec{N}$, there is a bijection β between $V(T)$ and $\mathcal{O}$ in such a way that there is an edge between two vertices v and w from T if and only if there is a unique separation s in N which distinguishes $\beta(v)$ and $\beta(w)$. Moreover, in that case $\alpha(v, w) = \vec{s}$, where $\vec{s}$ is the orientation of s contained in $\beta(w)$.*

Proof. For every orientation $O \in \mathcal{O}$ consider the orientation of $E(T)$ given by orienting $e = vw$ as (v, w) if and only if $\alpha(v, w) \in O$. Since O is a consistent orientation and (T, α) is order-respecting, no two edges of T can be oriented away from each other, thus this orientation has a unique sink v which we say is the node of T *containing* O. Moreover, since any two orientations in $\mathcal{O}$ are distinguished from each other by some separation in N and the image of α is all of $\vec{N}$, these sinks need to be distinct for distinct orientations from $\mathcal{O}$. If on the other hand there would be some node t of T which is not the sink for any orientation in $\mathcal{O}$, we could pick some edge $e = vt$ incident with t and delete the separation $\alpha(v, t)$ from N. Since any two orientations in $\mathcal{O}$ are contained in distinct nodes of T, this set would then contradict the $\subseteq$-minimality of N. Thus, let us define β by mapping a node $v \in T$ to the unique orientation in $\mathcal{P}$ contained in that node.

Now given any two adjacent nodes t, t' of T, the orientations of $E(T)$ induced by $\beta(t)$ and $\beta(t')$ are the same except for the edge tt'. Thus, the unoriented separation corresponding to $\alpha(t, t')$ is the only separation in N which distinguishes

$\beta(t)$ and $\beta(t')$, and clearly $\alpha(t,t') \in \beta(t')$. On the other hand, if two nodes t, t' of T are not adjacent then all the separations corresponding to $\alpha(e)$ for an edge e on the unique path between t and t' distinguish $\beta(t)$ and $\beta(t')$, in particular there are at least two such separations. □

In particular, the Lemmas 2.6.3 and 2.6.4 together imply that we can construct an S-tree from such a $\subseteq$-minimal nested set as follows:

Corollary 2.6.5. *Let $\mathcal{O}$ be a set of consistent orientations of a finite separation system S, let N be a $\subseteq$-minimal tree set which distinguishes any pair of orientations from $\mathcal{O}$ from another. Let T be the graph with vertex set $\mathcal{O}$ in which we add an edge to $E(T)$ between O and O' if and only if there is a unique separation in N which distinguishes O from O'. Let $\alpha : \vec{E}(T) \to \vec{N}$ map the directed edge from O to O' to the orientation of the unique separation in N distinguishing O and O' which is contained in O'. Then (T, α) is an order-respecting S-tree with image N.*

2.7 Tree-of-tangles theorems

Prior to this thesis, the two most general versions of a tree-of-tangles theorem were the following from [27] and [26]:

Theorem 1.1.2 (Canonical tree-of-tangles theorem for separation universes [27, Theorem 3.6])**.** *Let $U = (\vec{U}, \leqslant, ^*, \vee, \wedge, |\cdot|)$ 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.

Theorem 1.1.3 ([26, Theorem 6])**.** *Let $\vec{S}$ be a structurally submodular separation system and $\mathcal{P}$ a set of profiles of S. Then S contains a nested set that distinguishes $\mathcal{P}$.*

While the second of these two theorems removed the requirement of the existence of a submodular order function, the first of these statements gives a stronger result because of the resulting set distinguishing the profiles efficiently, and being canonical, i.e. invariant under automorphisms. A large part of Chapter 4 is devoted to obtain unified and more general versions of these two theorems, which allows one to combine the fewer assumptions of Theorem 1.1.3 with the stronger result of Theorem 1.1.2.

2.8 Examples

Let us now see some examples of these separation systems. These are mostly taken from [20, 21]. A typical example of a separation system are the *separations* *of a graph.* Given a (finite or infinite) graph G, an *(oriented) separation* of G is a pair (A, B) of vertex sets such that $A \cup B = V(G)$ and there is no edge from $A \setminus B$ to $B \setminus A$. The set $\vec{S}(G)$ of all these separations is a universe of separations when equipped with the involution $(A, B)^* = (B, A)$ and the partial order of declaring $(A, B) \leqslant (C, D)$ if and only if $A \subseteq C$ and $D \subseteq B$. Then $(A, B) \vee (C, D) = (A \cup C, B \cap D)$ and $(A, B) \wedge (C, D) = (A \cap C, B \cup D)$. We say that the order of a separation of G is $|A \cap B| \in \mathbb{N} \cup \{\infty\}$. If we consider the set $\vec{S}_{\aleph_0}(G)$ of all those separations of G of finite order, then this definition gives us a submodular order function on $\vec{S}_{\aleph_0}(G)$, thus $\vec{S}_{\aleph_0}(G)$ equipped with this order function is a submodular universe of separations. By a *(k-)profile in G* we shall mean a (k-)profile in $\vec{S}_{\aleph_0}(G)$.

In graphs, one historically does not consider k-profiles, but the slightly more restrictive notion of k-tangles. Given an integer k, a *k-tangle* of a graph is an orientation τ of $\vec{S}_k$ with the *tangle property:*

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

Note that every k-tangle τ of a graph needs to be a k-profile: if τ is not consistent, say because $(A, B), (C, D) \in \tau$ but $(B, A) \leqslant (C, D)$ then the pair $(A, B), (C, D)$ violates (T) as $G[A] \cup G[C] = G$. Similarly, τ cannot violate (P), as if τ contains $(A, B), (C, D)$ and $(B \cap D, A \cup C)$, then this violates (T) as $G[A] \cup G[C] \cup G[B \cap D] = G$. While every k-tangle is a regular k-profile, [27] gave an example of a k-profile in a graph that is not a k-tangle ([27, Example 7]).

A k-tangle of G is a *maximal* tangle of G if it is not the subset of some l-tangle of G for some $l > k$.

The notion of profiles originates in that of k-tangles of a graph. Likewise, also the notion of an S-tree has its origin in an important notion in graphs which is, as for example shown in [68], closely related to these tangles: the notion of a *tree-decomposition* of a graph G. A tree-decomposition is a pair $(T, \mathcal{V})$ of a tree T and a family $\mathcal{V} = (V_t)_{t \in T}$ of vertex sets $V_t \subseteq V(G)$ which satisfies the following three properties (see also [20]):

(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 that $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$.

The edges xy of the tree T of a tree-decomposition naturally induce a separation of the underlying graph: let T_x and T_y be the components of $T - xy$ containing x or y, respectively, and consider the separation

$$\left(\bigcup_{z \in T_x} V_x, \bigcup_{z \in T_y} V_y \right).$$

Consequently, a tree-decomposition corresponds to an S-tree by mapping every edge of T to the separation of G induced by that edge (see also [20, §12.5]). The *width* of a tree-decomposition is the largest size of a bag V_t minus 1. The *adhesion* of the tree-decomposition is the largest size of one of the sets $V_t \cap V_{t'}$, for distinct nodes t and t' of T.

The relation between k-tangles and tree-decompositions of low width (say less than l) lies in two important properties. On the one hand there is a 'duality': a graph has a tree-decomposition of low width only if the graph does not contain a k-tangle for a large k. While for tangles this relation is not 1:1 (in that $k = l$), the parameters depend on each other only by a linear factor (i.e. if a graph has a tree-decomposition of width less than l, the graph does not contain an $(l+1)$-tangle, and conversely, if the graph does not contain a tree-decomposition of width less than l, then the graph contains a $\lceil \frac{l}{3} \rceil$-tangle).

On the other hand, tree-decompositions can be used to distinguish distinct tangles. Let us say that a tree-decomposition $(T, \mathcal{V})$ of 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. The classic tree-of-tangles theorem by Robertson and Seymour can then be phrased as follows:

Theorem 1.1.1 ([68])**.** *Every graph has a tree-decomposition displaying its maximal tangles.*

The separations of a graph can be naturally generalized to another important class of separation systems, the *set separations* of a given set V. We say that, given $A, B \subseteq V$, that (A, B) is an (oriented) set separation of V if and only if $A \cup B = V$. For set separations, and also for separations of a graph, we usually write $\{A, B\}$ to denote the unoriented separation $\{(A, B), (B, A)\}$. Like the separations of a graph, these set separations form a universe of separations equipped with the partial order of defining $(A, B) \leqslant (C, D)$ if and only if $A \subseteq C$ and $D \subseteq B$. Moreover, if restricted to all the separations (A, B) with $A \cap B$ finite, there is a natural order function on the set separations of a set V given by $|(A, B)| = |A \cap B|$. A *universe of set separations* shall be any universe of these separations of a set with the above partial order. If we do not explicitly specify otherwise, we will implicitly assume that this universe comes with the order function mentioned above.

A subclass of these set separations is given by the *bipartition universe* $\mathcal{B}(V)$ consisting of only the bipartitions (A, B), that is only those separations (A, B) where $B = V \setminus A$. Note that $\mathcal{B}(V)$ as a lattice corresponds exactly to the subset lattice on V via the map $(A, B) \mapsto A$, since $(A, B), (C, D) \in \vec{\mathcal{B}}(V)$ satisfy $(A, B) \leqslant (C, D)$ precisely if $A \subseteq C$.

Chapter 3

Witnessing dense structures

We start the presentation of the results of this dissertation with results concerning concrete witnesses for the existence of dense structures both in graphs and abstract separation systems. The leading task in this chapter is to find concrete structures witnessing the existence of a 'highly connected substructure', for various notions of such structures. We start in Section 3.1 with tangles in graphs. We show that every such tangle admits a (weighted) decider, that is there is a weight function on the vertices of the graph, such that every separation in the tangle is oriented towards the side of higher weight. This weight function thus is, in some sense, a concrete 'highly connected substructure' corresponding to that tangle. This result is joint work with Jakob Kneip and Maximilian Teegen and can also be found in [37].

In Section 3.2 we are then concerned with the question, whether such weighted deciders exist for tangles in other contexts. While we already in Section 3.1 given an example of a context in which not every tangle admits a weighted decider, we in Section 3.2 search for sufficient conditions on tangles in contexts other than graphs to admit a weighted decider. We find a sufficient condition for this property to hold. Moreover, we also again consider tangles in graphs and ask under which condition we can actually find a decider *set* for such tangles, that is under which conditions does there exist a vertex set X in our graph, such that every separation (A, B) in our tangle satisfies $|A \cap X| < |B \cap X|$. These results are joint work with Reinhard Diestel and Raphael Jacobs and can also be found in [25].

After that, in Section 3.3, we consider another possible way to witness the existence of a tangle: via a tangle of a 'dual' separation system. We will describe a concrete setup of separation systems which are dual in some sense. In that setup we can then show that, a tangle of one of the two separation systems involved naturally defines an orientation of the other system which will, restricted to a slightly lower order, also be a tangle. Moreover, we will consider some variations of this idea. This joint work with Reinhard Diestel, Joshua Erde and Maximilian Teegen can also be found in [24].

Finally, in Section 3.4 we are concerned with another type of dense structures in graphs. Weißauer [76] suggested the notion of *agile sets* in graphs and asked whether the existence of large agile subsets of a graph can, quantitatively, be characterized via the existence of concrete minors of the graph. We are able to give such a characterization, and also consider some generalizations of these agile sets. These results are joint work with Maximilian Teegen and partly Jakob Kneip (see Appendix D) and are, at the time of writing, not published elsewhere.

3.1 Tangles are Decided by Weighted Vertex Sets

3.1.1 Introduction

A concrete example of a tangle in a graph is the following: if a graph 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 'defines τ by majority vote'. It results from this idea that every tangle is defined by 'majority vote', that we think, given some tangle τ in G containing a separation (A, B), of A and B as the 'small' and the 'big' side of (A, B) in that tangle, respectively. The main result of this section will make this intuition concrete.

Consequently, in [27] Diestel raised the question whether indeed 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 3.1.1. 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 my Master's thesis [35], where I showed that such a set X always exists if G is $(k-1)$-connected and has at least $4(k-1)$ vertices. However, this 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 section of this thesis, 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 so 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.

In Section 3.1.2 we will formulate and prove our main theorem asserting that tangles of graphs (and of hypergraphs) always admit such a weight function. Following that we show in Section 3.1.3 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.

For settings beyond graphs it is known that the analogue of Diestel's question may be false. For instance, Geelen [47] pointed out that there are matroid tangles which cannot be decided by majority vote, not even when considering a fractional version of the problem. For edge-tangles as analysed in Section 3.1.3 the fractional version of Problem 3.1.1 is true for graphs but may fail for hypergraphs. We demonstrate the latter with a counterexample which, though discovered independently, is conceptually similar to Geelen's example in the matroid setting.

3.1.2 Weighted deciders

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: since the separations of a graph form a separation system, there is a partial order on them. Hence, we can 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.

Recall that for a graph G the natural partial order on the separations of G is given by letting $(A, B) \leqslant (C, D)$ if and only if $A \subseteq C$ and $B \supseteq D$ and that a k-tangle of a graph is an orientation τ of all separations of order $< k$ with the property that τ does not contain any three separations $(A_1, B_1), (A_2, B_2), (A_3, B_3)$ so that $G[A_1] \cup G[A_2] \cup G[A_3] = G$. 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 the 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 3.1.2. *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 (A, B) and (C, D) 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 is modular, that is,

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

With the above inequalities this implies 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 3.1.3 (Tucker's Theorem [73])**.** *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.

Proof of 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 $w\colon V \to \mathbb{R}_{\geqslant 0}$ so 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

$$m_{ij} \coloneqq |B_i \cap (A_j \cap B_j)| - |A_i \cap (A_j \cap B_j)|$$

for $i, j \leqslant n$. Let M be the matrix $\{m_{ij}\}_{i,j\leqslant n}$. Observe that, by Lemma 3.1.2, 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

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

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 3.1.3 to K. We define a weight function $w\colon V \to \mathbb{R}$ by

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

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

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

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 $\varepsilon > 0$ such 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 this section with the remark that Theorem 1 and its proof extend to tangles in hypergraphs without any changes. Even more generally, the following version of Theorem 1, which is formulated in terms of profiles of set separations, can be established with exactly the same proof as well:

Theorem 3.1.4. *Let $\vec{U}$ be a universe of set separations of a finite ground-set V with the order function $|(A, B)| := |A \cap B|$. Then, for any regular k-profile P in $\vec{U}$, 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 $w(A) < w(B)$.*

Recall that a set separation of some ground-set V is a pair (A, B) of subsets of V with $A \cup B = V$ and that a set $\vec{U}$ of such separations is a universe if $\vec{U}$ contains (B, A) and $(A \cup C\,,\, B \cap D)$ for all (A, B) and (C, D) in $\vec{U}$. Like for the separations of a graph, a partial order on the set separations of V is given by letting $(A, B) \leqslant (C, D)$ if $A \subset C$ and $B \supseteq D$.

Recall furthermore that, for an integer k, a regular k-profile in $\vec{U}$ is a set P consisting of exactly one of (A, B) and (B, A) for every (A, B) in $\vec{U}$ of order $|A \cap B| < k$, with the additional property that there are no (A, B) and (C, D) in P for which $(B, A) \leqslant (C, D)$ or such that P contains $(B \cap D\,,\, A \cup C)$.

Observe that if $G = (V, E)$ is a (hyper-)graph then the set $\vec{U}$ of all separations of G is such a universe. Moreover, every k-tangle τ of G is also a regular k-profile of $\vec{U}$. Therefore, Theorem 3.1.4 indeed applies to tangles in graphs and hypergraphs as well.

Theorem 3.1.4 holds with the same proof as Theorem 1, since Lemma 3.1.2 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.

3.1.3 Edge-tangles

A related object of study (cf. [29, 63]) 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. 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 2. *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 *pairwise weighting* to distinguish it from the weight function of a decider. The *order* of a bipartition (A, B) is defined as $|(A, B)| := \sum_{(u,v)\in A\times B} e(u, v)$. Note that this function is submodular in the sense that for all bipartitions (A, B) and (C, D) we have

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

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 does not contain $(V, \emptyset)$ and has no subset of the form $\{(A, B), (C, D), (B \cap D, A \cup C)\}$.

Observe that every k-edge-tangle of a (multi-)graph $G = (V, E)$ is also a k-profile in (V, e), where e is the multiplicity of the edges of G. Therefore, the following theorem directly implies Theorem 2:

Theorem 3.1.5. *Let (V, e) be a pairwise 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 on the edges by the same principles as in Theorem 1 and to then transform it into the weighted vertex decider w. So let us first show an analogue of Lemma 3.1.2 for pairwise 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 3.1.2:

Lemma 3.1.6. *For every r-profile τ in a pairwise 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 3.1.5:

Proof of Theorem 3.1.5. Like in the proof of Theorem 1, it suffices to find a suitable real-valued weight function $w\colon V \to \mathbb{R}_{\geqslant 0}$ since V is finite. We will begin by finding a weight function $\overline{w}\colon V^2 \to \mathbb{R}_{\geqslant 0}$ on the pairs in V so that we have $\overline{w}(A) \leqslant \overline{w}(B)$ for all $(A, B) \in \tau$, where $\overline{w}(A) = \sum_{(u,v) \in A^2} \overline{w}(u, v)$, and with this inequality being strict for all but at most one of the maximal elements of τ. We will subsequently use this $\overline{w}$ to construct the desired weight function $w\colon V \to \mathbb{R}_{\geqslant 0}$.

Enumerate the maximal elements of τ as $(A_1, B_1), \dots, (A_n, B_n)$. Just as in Theorem 1 it suffices to find a weight function which decides these maximal elements. For every two maximal elements (A_i, B_i) and (A_j, B_j) let

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

Let M be the matrix $\{m_{ij}\}_{i,j \leqslant n}$. Observe that, by Lemma 3.1.6, $M + M^T$ 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 $x \in \mathbb{R}^n_{\geqslant 0}$ so 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

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

Note that $\overline{w}$ is symmetric. For the same reason as in Theorem 1, by choice of x, this function $\overline{w}$ decides all but at most one of the (A_i, B_i) correctly in the sense that $\overline{w}(A_i) \leqslant \overline{w}(B_i)$ for all $i = 1, \dots, n$ with at most one inequality not being strict.

It remains to turn $\overline{w}$ into a weight function on V rather than on V^2, and to verify that it has the desired properties. Define $w\colon V \to \mathbb{R}_{\geqslant 0}$ as

$$w(v) \coloneqq \sum_{u \in V} \overline{w}(u, v)\,.$$

Then, for each $i = 1, \dots, n$, we find that

$$\begin{aligned}
w(B_i) - w(A_i) &= \sum_{u \in B_i} \sum_{v \in V} \overline{w}(u,v) - \sum_{u \in A_i} \sum_{v \in V} \overline{w}(u,v) \\
&= \sum_{(u,v) \in B_i^2} \overline{w}(u,v) - \sum_{(u,v) \in A_i^2} \overline{w}(u,v) \\
&= \sum_{(u,v) \in B_i^2} \sum_{\substack{j: \\ (u,v) \in (A_j \times B_j) \cup (B_j \times A_j)}} x_j \cdot e(u,v) \quad - \sum_{(u,v) \in A_i^2} \sum_{\substack{j: \\ (u,v) \in (A_j \times B_j) \cup (B_j \times A_j)}} x_j \cdot e(u,v) \\
&= 2 \sum_{j=1}^{n} \left(\sum_{(u,v) \in B_i^2 \cap (A_j \times B_j)} x_j \cdot e(u,v) \; - \sum_{(u,v) \in A_i^2 \cap (A_j \times B_j)} x_j \cdot e(u,v) \right) \\
&= 2(Mx)_i.
\end{aligned}$$

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. □

Remarkably, and in contrast to Theorem 1, Theorem 2 does not in fact extend to hypergraphs. To see this, let us recall the relevant definitions, which extend naturally to hypergraphs.

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 H is a bipartition (A, B) of V and 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 m, a *m-edge-tangle* of H is a set τ consisting of exactly one (A, B) or (B, A) for every cut (A, B) of order $< m$, 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 m hyperedges of H.

A weighted decider for some m-edge-tangle τ of a hypergraph $H = (V, E)$ then is a function $w \colon V \to \mathbb{N}$ such that a cut (A, B) of H of order $< m$ lies in τ if and only if $w(A) < w(B)$.

Theorem 2 thus asserts that if H is just a (multi-)graph, i.e. if every hyperedge in E has size 2, then every m-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 with hyperedges of size $\geqslant 3$.

Example 3.1.7. For some natural number $m \geqslant 6$ let k be an integer with $3 \leqslant k \leqslant \frac{m}{2}$. Let V be the set of all k-element subsets of $[m] = \{1, \dots, m\}$. Let the set E of hyperedges consist of, for each $i \in [m]$, the set of all $v \in V$ that contain i. Note that each of these m many hyperedges of H has size $\binom{m-1}{k-1}$, making H a uniform k-regular hypergraph.

Theorem 3.1.8. *Let H be as in Example 3.1.7. Then H has a m-edge-tangle with no weighted decider.*

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

We can therefore define

$$\tau := \{(A, B) \in S_m \;:\; \cup A \subsetneq [m]\} \,.$$

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

To see that τ is a m-edge-tangle we note that by the above observation τ contains exactly one of (A, B) or (B, A) for every cut $(A, B) \in S_m$. 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 $[m] \setminus \cup A_1$, from $[m] \setminus \cup A_2$, and from $[m] \setminus \cup A_3$ lies in $B_1 \cap B_2 \cap B_3$, which is hence non-empty since such a $v \in V$ exists by $k \geqslant 3$. Finally for each $(A, B) \in \tau$ the set B is incident with each hyperedge of H since $\cup B = [m]$. Thus, τ is indeed a m-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 $i \in [m]$ 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 $i \in [m]$. We therefore have

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

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

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

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 $k \leqslant \frac{m}{2}$ that

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

Therefore, there can be no weighted decider for τ. □

A construction analogous to Example 3.1.7 was found independently by Geelen [47] 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 3.1.7 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 $m = 2k$ the tangle described in Theorem 3.1.8 has no weighted decider with real-valued and possibly negative weights either, with the same proof.

3.2 Deciders for tangles of set partitions

3.2.1 Introduction

While Theorem 3.1.4 and Theorem 2 provide the existence of a weighted decider for tangles and profiles in a lot of different contexts, there still exist contexts in which they do not guarantee this. For example, we cannot guarantee the existence of a weighted decider for edge-tangles in hypergraphs, as Theorem 3.1.8 shows that not each such tangle has a weighted decider. Therefore, it seems natural to ask for sufficient additional conditions a tangle or profile may satisfy and which do guarantee the existence of a weighted decider even if we are not in the context of Theorem 3.1.4 and Theorem 2. We searched for such conditions for quite a while, until we came up with the following: let $\vec{S}$ be any abstract separation system in some universe $\vec{U}$ of separations. Now given any $k \in \mathbb{N}$, we say that an orientation τ of S is *k-resilient* if no set of $\leqslant k$ elements of τ has a co-small supremum in $\vec{U}$.

Both tangles and edge-tangles of graphs as defined in Section 3.1.3 are 3-resilient. The example of a hypergraph tangle constructed in Theorem 3.1.8 with parameters k and l is l-resilient, but its ground set contains $\binom{k}{l}$ many elements. On the other hand, every *principal* profile in a universe of bipartitions, that is a profile consisting of all those bipartitions (A, B) whose big side B contains some fixed element x, is *infinitely resilient* in that it is k-resilient for every $k \in \mathbb{N}$. Note that $\{x\}$ is a decider set for this profile.

These examples seem to suggest that profiles of bipartitions, or more generally profiles of set separations, that are k-resilient for large k are more likely to have decider sets. We can indeed prove such a fact, with an interesting additional twist: 'large' has to be measured not in terms of $|V|$ or $|S|$, but relative to the number of those elements of the profiles which are maximal in the profile with respect to the partial order of oriented separations. This as such is unsurprising: in a k-resilient profile with at most k maximal elements, the intersection of all their big sides is non-empty, and is clearly a decider set for this profile.

The exact statement we can show is the following:

Theorem 3. *Let $\vec{U}$ be the universe of all set separations of some finite set V, and let τ be an orientation of some set $S \subseteq U$ of separations. Let m be the number of maximal elements of τ. If τ is k-resilient for some $k > \frac{m}{2}$, then τ has a decider.*

Moreover, we will also see that there are, for every $k \leqslant \frac{m}{2}$, examples of tangles which are k-resilient but do not have a decider.

The second part of this section of this thesis deals with the existence of decider *sets* for profile of set separations, that is in the universe U of set separations of a finite set V equipped with the submodular order function assigning order $|A \cap B|$ to the separation $s = \{A, B\} \in U$.

As we have shown in Theorem 3.1.4, each k-profile of this U has a weighted decider. But what if we are instead interested in a decider set? We will provide a partial answer to this question in Section 3.2.4 by showing that, if a k-profile extends to a profile twice its order, then it has a decider *set*, i.e. we will be able to show the following:

Theorem 4. *Let $\vec{U}$ be a universe of set separations of some finite set V equipped with the order function $|(A, B)| = |A \cap B|$ and let $k \in \mathbb{N}$. If τ' is a k-profile in $\vec{U}$*

which extends to a regular $2k$-profile τ in $\vec{U}$, then τ' has a decider set $X \subseteq V$ of size $|X| \geqslant 2k$.

After giving the basis definitions about deciders used in this section of this thesis in Section 3.2.2, we start in Section 3.2.3 with proving Theorem 3. After that, in Section 3.2.4 we prove Theorem 4.

3.2.2 Weight Functions and Deciders

Let us define a *weight function* on a finite set V as any map w from V to $\mathbb{R}_{\geqslant 0}$. For any subset $U \subseteq V$ we write $w(U) = \sum_{v \in U} w(v)$. Note that

$$w(B) - w(A) = w(B \setminus A) - w(A \setminus B)$$

for every set separation (A, B) of V and every weight function w on V, a fact we shall use freely throughout. A weight function w on V with values in $\{0, 1\}$ can be equivalently formulated as an indicator function of the set $X = X_w = w^{-1}(1)$ in that $w(A) = |X \cap A|$ for every $A \subseteq V$. For such a weight function we shall transfer all the wording around weight functions equivalently to the set X.

Let w be a weight function on a set V. We say that w *decides* a set separation $s = \{A, B\}$ of V if there is an orientation $\vec{s} = (A, B)$ of s such that $w(A) < w(B)$; we shall also say that w *decides s as $\vec{s}$*. If S is a separation system which consists of set separations of V, then w *decides S* if it decides each of its elements. In particular, the fact that w decides S yields an orientation τ of S by orienting a separation $s \in S$ as $\vec{s}$ if w decides s as $\vec{s}$.

Now given an orientation τ of S we say that a weight function w *decides S as τ* if w decides every $s \in S$ as $\vec{s} \in \tau$, i.e. if $w(A) < w(B)$ for every $(A, B) \in \tau$. Such a weight function w is called a *decider (function)* for τ. If the weight function w has values in $\{0, 1\}$, i.e. it comes from the set $X = w^{-1}(1)$, then X (and also w) is a *decider set* for τ.

A weight function w *witnesses* a separation (A, B) if $w(A) < w(B)$. Thus, a decider w for some previously given orientation τ of S witnesses all separations in τ. We therefore say that w *witnesses* τ; if a decider witnessing τ exists, then we say that τ *has* a decider (or a witness).

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

$$\lambda = \frac{K}{\min_{(A,B) \in \tau}(w(B) - w(A))}.$$

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

Also, as in Section 3.1 it is enough to consider the maximal separation in an orientation τ if one is concerned with the question of whether a given function w decides S as τ. More precisely, we have the following easy observation:

Observation 3.2.1. Let $w : V \to \mathbb{R}_{\geqslant 0}$ be a weight function on a set V. Let (A, B) and (C, D) be two set separations of V with $(C, D) \leqslant (A, B)$. If w witnesses (A, B), then it witnesses (C, D) as well.

Proof. Since $(C, D) \leqslant (A, B)$, we have $C \subseteq A$ and $D \supseteq B$. So $w(C) \leqslant w(A)$ and $w(D) \geqslant w(B)$, as w is a weight function. Now w witnesses (A, B) in that $w(A) < w(B)$. This immediately implies

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

so w witnesses (C, D) as well. □

3.2.3 Deciders and resilience

In this section we use the novel notion of *resilience* to prove a sufficient criterion for an orientation of some set of set separations to have a decider. We begin this section by giving all the definitions around the concept of resilience.

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

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

If for some orientation τ there exists a maximal $k \in \mathbb{N}$ such that τ is k-resilient, then we call k the *resilience* of τ.

Let us illustrate the concept of resilience with an example. Consider the 5-tangle τ of the $(n \times n)$-grid that has the entire vertex set of the grid as a decider set. Let us show that τ is $\Omega(n^2)$-resilient. Notice that every element (A, B) of τ satisfies $|A| \leqslant 10$; indeed, most satisfy $|A| \leqslant 5$. Since all co-small set separations of V are of the form (V, X) for some $X \subseteq V$, any set of separation in τ with a co-small supremum has at least $\frac{n^2}{10}$ elements.

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

Assume that τ is k-resilient, for some integer k, and write $\mathcal{M}$ for the set of all k-element subsets of $M = M(\tau)$. Then, for every $M' \in \mathcal{M}$, there exists an element $v_{M'}$ of our ground set V which is strictly on the big side of all separations in M'. It seems natural to construct a decider for M (and thus for τ by Observation 3.2.1) by combining all these local decider sets $\{v_{M'}\}$: we assign to each $v \in V$ as its weight the number of sets $M' \in \mathcal{M}$ with $v_{M'} = v$.

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

More precisely, we have the following theorem:

Theorem 3. *Let $\vec{U}$ be the universe of all set separations of some finite set V, and let τ be an orientation of some set $S \subseteq U$ of separations. Let m be the number of maximal elements of τ. If τ is k-resilient for some $k > \frac{m}{2}$, then τ has a decider.*

Proof. Let $M(\tau)$ be the set of maximal elements of τ. Given some $M' \subseteq M(\tau)$ of size k, we know that there exists, by the definition of k-resilience, an element $v_{M'} \in \bigcap_{(A,B)\in M'} B$. Let $v_{M'}$ be chosen as an arbitrary such element of $\bigcap_{(A,B)\in M'} B$

We define our weight function $w : V \to \mathbb{R}_{\geqslant 0}$ by defining $w(v)$, for $v \in V$, as the number of subsets M' of $M(\tau)$ of size k, for which $v = v_{M'}$ and claim that this function is a weighted decider for O.

Indeed, by Observation 3.2.1 it is enough to show that w is a decider for all separations in $M(\tau)$. However, given some $(A,B) \in M(\tau)$, there are $\binom{m-1}{k-1}$ many distinct subsets $M' \subseteq M(\tau)$ of size k which contain (A,B) and $\binom{m-1}{k}$ many such subsets $M' \subseteq M(\tau)$ of size k which do not contain (A,B). Since $k > \frac{m}{2}$ we observe that $\binom{m-1}{k-1} > \binom{m-1}{k}$. Thus,

$$\begin{aligned} w(A) = \sum_{a\in A} w(a) &\leqslant |\{M' \subseteq M(\tau) \,:\, |M'| = k, (A,B) \notin M'\}| \leqslant \binom{m-1}{k} \\ &< \binom{m-1}{k-1} = |\{M' \subseteq M(\tau) \,:\, |M'| = k, (A,B) \in M'\}| \leqslant \sum_{b\in B} w(b) = w(B) \end{aligned}$$

Thus, w decides every separation $(A,B) \in M(\tau)$ as (A,B) and thus is, by Observation 3.2.1, a decider for τ. ☐

As it turns out, this bound of $k > \frac{m}{2}$ in Theorem 3 is optimal in the following sense:

Proposition 3.2.2. *For every $k, m \in \mathbb{N}$ with $3 \leqslant k \leqslant \frac{m}{2}$, there exists an edge-tangle $\tau_{m,k}$ of a hypergraph with m maximal elements that is k-resilient, but which has no decider.*

Proof. Consider the uniform k-regular hypergraph H from Example 3.1.7. Let us denote as $\tau_{m,k}$ the m-tangle τ constructed in Theorem 3.1.8 for this hypergraph H.

Let us denote, for every $i \in [m]$, as V_i the set of all those vertices of H which correspond to a k-element subset of $[m]$ containing i. Then the maximal separations in $\tau_{m,k}$ are, by construction, all the bipartitions $\vec{s}_i = (V \smallsetminus V_i, V_i)$.

We only need show that $\tau_{m,k}$ is indeed k-resilient, but this is immediate from the construction: by Observation 3.2.1, it is enough to consider an arbitrary collection $\vec{s}_{i_1}, \ldots, \vec{s}_{i_k}$ of k separations in the set M of maximal separations in $\tau_{m,k}$. But then there is a vertex of H corresponding to the k-element subset $\{i_1, \ldots, i_k\}$

of $[m]$, and this vertex is on the big side V_{i_j} of $\vec{s}_{i_j}$ for all $j \in [j]$ by construction. □

Now another approach we might take in order to guarantee the existence of a decider for our given profile τ of set separations is to make further requirements that τ should satisfy. Let us for the moment only consider the case of bipartitions. We might increase our requirements on τ by requiring that, for some fixed integer ℓ we have that $|B_1 \cap B_2 \cap B_3| \geqslant \ell$ for any three $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$. So let us say that an orientation τ of bipartitions is an $\mathcal{F}_\ell$-tangle if τ does not contain any three separations $(A_1, B_1), (A_2, B_2), (A_3, B_3)$ so that $|B_1 \cap B_2 \cap B_3| < \ell$. Now for any fixed ℓ, the construction from Example 3.1.7 and Theorem 3.1.8 can easily be modified to find an $\mathcal{F}_\ell$-tangle which still does not admit a weighted decider, but is k-resilient for our favourite $k \leqslant \frac{m}{2}$: we can simply take ℓ disjoint copies of every vertex of H and let any hyperedge of H contain all copies of the vertices contained in that edge.

However, maybe we can guarantee the existence of a decider by requiring that ℓ is not only a large constant, but large in terms of $|V|$, e.g. of size at least $\epsilon |V|$ for some constant $\epsilon > 0$. The following proposition shows that there exists a sharp lower bound for those $\epsilon > 0$ for which $\ell \geqslant \epsilon |V|$ guarantees the existence of decider:

Proposition 3.2.3. *Let V be a (finite) set, and let $0 < \epsilon < 1$. If $\epsilon \geqslant \frac{1}{8}$, then every $\mathcal{F}_{\epsilon n}$-tangle τ of some set S of bipartitions of V has a weighted decider; if $\epsilon > \frac{1}{8}$, then τ even has a decider set.*

Conversely, for every $\epsilon < \frac{1}{8}$ we find some $n \in \mathbb{N}$, so that for corresponding choices of m and k the tangle $\tau_{m,k}$ from Proposition 3.2.2 forms an example of an $\mathcal{F}_{\epsilon n}$-tangle which orient bipartitions of a ground set V of size n and does not have a decider.

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

It turns out that if $\ell > \frac{|V|}{8}$, then $B_1 \cap B_2$ needs to be the desired decider set: otherwise, there exists another separation $(A_3, B_3) \in \tau$ such that $|(B_1 \cap B_2) \cap B_3| \leqslant |(B_1 \cap B_2) \cap A_3|$. This implies $|B_1 \cap B_2 \cap B_3| \leqslant \frac{|V|}{8}$ which contradicts the fact that τ is an $\mathcal{F}_\ell$-tangle.

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

Suppose there are two separations $(A, B), (C, D) \in \tau'$ with $|B \cap C| > |B \cap D|$. Then this yields $|B \cap D| < \frac{|V|}{4}$ which in turn implies the existence of a decider set with the same arguments as above. Consequently, for every two separations $(A, B), (C, D) \in \tau'$, we have $|B \cap C| \leqslant |B \cap D|$.

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

For the second part of the proposition, let us consider the tangle $\tau_{m,k}$ from Proposition 3.2.2 for some $m \geqslant 2k \geqslant 6$. Then, for any three maximal separations in $\tau_{m,k}$, their intersection contains exactly $\binom{m-3}{k-3}$ elements of the constructed ground set V. In particular, $\tau_{m,k}$ is an $\mathcal{F}_\ell$-tangle for all $\ell < \binom{m-3}{k-3}$. Now recall that the size of the ground set V is $|V| = \binom{m}{k}$. Thus, if we set $k = \frac{m}{2}$, then

$$\lim_{m\to\infty} \frac{\binom{m-3}{k-3}}{\binom{m}{k}} = \frac{1}{8}.$$

So, by Proposition 3.2.2, we find for any $\epsilon < \frac{1}{8}$ some $n \in \mathbb{N}$ and integers m, k so that the tangle $\tau_{m,k}$ witness that there exists an $\mathcal{F}_{\epsilon n}$-tangle on a ground set of n elements without a weighted decider. □

3.2.4 Deciders for extendable tangles

Let $\vec{U}$ be a universe of set separations of a ground set V which is equipped with a submodular order function, and let τ be an orientation of S_k for some $k \in \mathbb{N}$. In Section 3.2.3 we analysed different properties of τ which ensure the existence of a weighted decider for τ. All the notions we considered there may be viewed as additional requirements τ need to satisfy in order for us to guarantee the existence of a decider for τ.

But how can we guarantee the existence of a decider for a profile τ if we do not want to impose additional requirements on τ in the above sense? We know that there exist profiles, and even tangles, without deciders (see e.g., Proposition 3.2.2). So instead of looking for a decider for τ itself, we may try to find a decider for some subset τ' of the separations of τ. Ideally, we can do so in such a way that this decider for the subset is still, in some sense, related to the original profile τ.

In the presence of an order function, one natural such subset of a k-profile τ, say, consists of all separations of order less than some $k' < k$. In other words, we would like to obtain, given a k-profile τ, a decider for the k'-profile $\tau' \subseteq \tau$. One way in which we could try to achieve this consists in proving the following: if a profile τ' extends to some profile τ of higher order in $\vec{U}$, then τ' has a decider. Here, we say that an orientation τ' of a separation system $S' \subseteq S$ *extends* to an orientation τ of S if $\tau \cap \vec{S'} = \tau'$. In this case, we may view the decider w for τ' as an approximation of a decider for its extension τ – although w will in general not decide all the separations in τ. The profiles constructed in Proposition 3.2.2, for example, do not have a decider; but if we consider only those separations of order at most $\frac{m}{2}$ in this example, then they even have a decider set: the whole ground set V decides all separations of order at most $\frac{m}{2}$ as in $\tau_{m,k}$.

This leads us to the question whether profiles which extend to profiles of twice their order do always have decider sets.

In this section we show that k-profiles which extend to regular $2k$-profiles do indeed have decider sets – if we work in a universe $\vec{U}$ of set separations equipped with the order function $|(A, B)| = |A \cap B|$. Recall that in this setting, we already know from Theorem 3.1.4 that there exists a weighted decider for any k-profile in $\vec{U}$. The following theorem strengthens this in that it shows the existence of an actual decider set for certain k-profiles:

Theorem 4. *Let $\vec{U}$ be a universe of set separations of some finite set V equipped with the order function $|(A, B)| = |A \cap B|$ and let $k \in \mathbb{N}$. If τ' is a k-profile in $\vec{U}$ which extends to a regular $2k$-profile τ in $\vec{U}$, then τ' has a decider set $X \subseteq V$ of size $|X| \geqslant 2k$.*

The proof of Theorem 4 will find a star $\sigma \subseteq \tau$ whose *interior* $\bigcap_{(A,B)\in\sigma} B$ is the desired decider set for τ'. Let us show first that the interior of any star in τ has size at least $2k$.

Lemma 3.2.4. *If τ is a regular $2k$-profile in $\vec{U}$ for some $k \in \mathbb{N}$ and $\sigma \subseteq \tau$ is a star, then the interior $X = \bigcap_{(A,B)\in\sigma} B$ of σ satisfies $|X| \geqslant 2k$.*

Proof. Suppose that this is not the case, and let $\sigma \subseteq \tau$ be a star such that its interior $X = \bigcap_{(A,B)\in\sigma} B$ has size $|X| < 2k$.

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

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

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

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

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

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

Proof of Theorem 4. Let σ be a star in τ with an interior $X = \bigcap_{(A,B)\in\sigma} B$ of smallest possible size. By Lemma 3.2.4 we have $|X| \geqslant 2k$. We claim that X is the desired decider set for τ'.

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

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

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

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

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

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

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

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

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

completing the proof. □

Our proof of Theorem 4 heavily relies on the assumption that the order function on $\vec{U}$ is given by $|(A, B)| = |A \cap B|$. We do not know whether a similar result holds for other or even all submodular order functions on such $\vec{U}$:

Problem 3.2.5. Let $\vec{U}$ be a universe of set separations of a finite set V, and suppose that $\vec{U}$ is equipped with some submodular order function. Is it true that, if τ' is a k-profile in $\vec{U}$ for some $k \in \mathbb{N}$ which extends to a regular $2k$-profile in $\vec{U}$, then τ has a decider set? What if we consider other universes of set separations like e.g. the universe of bipartitions of V?

3.3 Dual separation systems

Unlike in the previous sections, in this section we do not try to witness the existence of a tangle by some vertex set, instead we are dealing with a specific scenario in which a tangle might be witnessed by a tangle of a different, in some sense dual, separation system.

As a concrete example let us think about the following setup: suppose we are given a set of objects, and a list of properties such that each object may, or may not, have each of the properties. This can be visualized by a bipartite graph, where the vertices on the sides correspond to the set X of objects and the set Y of properties, respectively, and we draw an edge between an object and a property precisely if the object has that property. Now we might be interested in tangles on the set X of objects, say of the set $\vec{S}(X)$ of set separations or the set $\vec{\mathcal{B}}(X)$ of bipartitions of X. Clearly, we want to restrict ourselves to a subset of these separations, as there are no tangles of either $\vec{S}(X)$ or $\mathcal{B}(X)$, and one natural to do so is by considering a submodular order function. This order function now clearly should take into account how the edges between X and Y are distributed, as we would think of two objects as being rather similar if they agree for a lot of properties on whether the have, or fail to have, that specific property.

However, in this specific setup, we might also be interested in tangles of the set Y of *properties*, we might want to analyse which properties often occur together, and so we might look at tangles of a subset of $\vec{S}(Y)$ or $\vec{\mathcal{B}}(Y)$, the set separations or bipartitions of the set Y of properties, instead. This then results in a dual approach to the one above, where we want to consider an order function on separations of Y, which again should take into account how the edges between X and Y are distributed.

We would suspect that the two types of tangles – those in X and those in Y – are linked in some way, and it might be the case that a given tangle of separations of X, say, witness that there is some tangle of separations of Y, and the other way around.

In this section we are giving a formal setup in which this is indeed possible, and we will show how to do so.

The setup works with the sets $\vec{S}(X)$ and $\vec{S}(Y)$ of set separations of X and Y, however we can find a similar setup for bipartitions, and will show how to do so at the end of Section 3.3.3.

3.3.1 Tangles on the sides of a bipartite graph

Given a bipartite graph G with partition classes X and Y, let us 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 will 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 gives us a natural way to partition the vertices in Y.

So, given $(A,B) \in \vec{S}(X)$ we define the separation $(A,B)^{\triangleright} := (A^{\triangleright}_B, B^{\triangleright}_A) \in \vec{S}(Y)$ by letting

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

and

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

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

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

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

and

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

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 map ${}^{\triangleleft}(\cdot) \colon 2^{\vec{S}(Y)} \to 2^{\vec{S}(X)}$ which we call a *pull-back*, sending every $\tau \subseteq \vec{S}(Y)$ to

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

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 sets $\vec{S}_k(X)$ and $\vec{S}_k(Y)$ of separations 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 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

[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 $^{\triangleright}$ and $^{\triangleleft}$ 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.

vertex from the appropriate class. For example, in the motivational example from the introduction, the order of a separation (A, B) of X will be determined by how evenly this separation splits, for every property in Y, the set of objects sharing that property. The more balanced the split, the larger the contribution of this property to the order of the separation. Back to the general setup of a bipartite graph G with partition classes X and Y, the order of a separation (A, B) will be the lower, the more vertices from the opposite partition class have a clear 'preference' of one side or the other.

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

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

Here, the first term will be larger when $N(y)$ is more evenly split by (A, B).

The reason why we multiplied the term $\min\{|N(y) \cap A|, |N(y) \cap B|\}$ with a factor of two and added the extra term of $-|N(y) \cap A \cap B|$ is to adjust for double-counting: if x is contained in both A and B we would count all the edges incident with x in $\sum_{y \in Y} \min\{|N(y) \cap A|, |N(y) \cap B|\}$, so moving x out of A or B would not increase the order even if it made the separation more balanced. Now in order to prevent this and still keep the order of a separation integer valued, we added the factor 2 before $\min\{|N(y) \cap A|, |N(y) \cap B|\}$ as well as the extra term $-|N(y) \cap A \cap B|$. This then results in a neighbour x in $A \cap B$ is treated as lying half in A and half in B.

Similarly, we define $|\cdot|_Y : \vec{S}(Y) \to \mathbb{N}$ where

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

Note that these functions are symmetric and non-negative, as required of an order function for separation systems. Less obviously, they are submodular, so $\vec{S}(X)$ and $\vec{S}(Y)$ equipped with these functions are submodular universes.

Moreover, the function $|\cdot|_X$ attains its maximum value on the degenerate separation (X, X), and since orientations of all of $\vec{S}(X)$ are not enlightening, in particular no such orientation will be a tangle or a profile, we will in the following implicitly assume that any $\vec{S_k}(X)$ we consider does not contain the separation (X, X).

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 three lemmas at first reading, to remain with the flow of the narrative.

Lemma 3.3.1. *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

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

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 $N_Z := |N(y) \cap Z|$.

Take separations (A_1, B_1) and (A_2, B_2) 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 $|(A_i, B_i)|_y = 2N_{A'_i} + N_{Z_i}$.

We observe that

$$|(A_1 \cap A_2, B_1 \cup B_2)|_y = 2N_{A'_1 \cap A'_2} + N_{Z_1 \cap Z_2} + N_{Z_1 \cap A'_2} + N_{A'_1 \cap Z_2},$$

and

$$|(A_1 \cup A_2, B_1 \cap B_2)|_y = 2\min\{N_{A'_1 \cup A'_2}, N_{B'_1 \cap B'_2}\} + N_{Z_1 \cap Z_2} + N_{Z_1 \cap B'_2} + N_{B'_1 \cap Z_2}.$$

Summing these two, we get

$$\begin{aligned}
&|(A_1 \cap A_2, B_1 \cup B_2)|_y + |(A_1 \cup A_2, B_1 \cap B_2)|_y \\
&= 2N_{A'_1 \cap A'_2} + N_{Z_1 \cap Z_2} + N_{Z_1 \cap A'_2} + N_{A'_1 \cap Z_2} \\
&\quad + 2\min\{N_{A'_1 \cup A'_2}, N_{B'_1 \cap B'_2}\} + N_{Z_1 \cap Z_2} + N_{Z_1 \cap B'_2} + N_{B'_1 \cap Z_2} \\
&\leqslant 2N_{A'_1 \cap A'_2} + 2N_{A'_1 \cup A'_2} \\
&\quad + N_{Z_1 \cap Z_2} + N_{Z_1 \cap A'_2} + N_{Z_1 \cap B'_2} + N_{Z_1 \cap Z_2} + N_{A'_1 \cap Z_2} + N_{B'_1 \cap Z_2} \\
&= 2N_{A'_1} + 2N_{A'_2} + N_{Z_1} + N_{Z_2} \\
&= |(A_1, B_1)|_y + |(A_2, B_2)|_y.
\end{aligned}$$

Similarly,

$$|(A_1 \cap B_2, B_1 \cup A_2)|_y = 2\min\{N_{A'_1 \cap B'_2}, N_{B'_1 \cup A'_2}\} + N_{Z_1 \cap Z_2} + N_{Z_1 \cap B'_2} + N_{A'_1 \cap Z_2}.$$

and

$$|(A_1 \cup B_2, B_1 \cap A_2)|_y = 2\min\{N_{A'_1 \cup B'_2}, N_{B'_1 \cap A'_2}\} + N_{Z_1 \cap Z_2} + N_{Z_1 \cap A'_2} + N_{B'_1 \cap Z_2}.$$

Summing these two, we get

$$\begin{aligned}
&|(A_1 \cap B_2, B_1 \cup A_2)|_y + |(A_1 \cup B_2, B_1 \cap A_2)|_y \\
&= 2\min\{N_{A'_1 \cap B'_2}, N_{B'_1 \cup A'_2}\} + N_{Z_1 \cap Z_2} + N_{Z_1 \cap B'_2} + N_{A'_1 \cap Z_2} \\
&\quad + 2\min\{N_{A'_1 \cup B'_2}, N_{B'_1 \cap A'_2}\} + N_{Z_1 \cap Z_2} + N_{Z_1 \cap A'_2} + N_{B'_1 \cap Z_2} \\
&\leqslant 2N_{A'_1 \cap B'_2} + 2N_{A'_2 \cap B'_1} \\
&\quad + N_{Z_1 \cap Z_2} + N_{Z_1 \cap A'_2} + N_{Z_1 \cap B'_2} + N_{Z_1 \cap Z_2} + N_{A'_1 \cap Z_2} + N_{B'_1 \cap Z_2} \\
&\leqslant 2N_{A'_1} + 2N_{A'_2} + N_{Z_1} + N_{Z_2} \\
&= |(A_1, B_1)|_y + |(A_2, B_2)|_y.
\end{aligned}$$

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

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

$$|(A, B)|_X = 2|E(A^{\triangleright}_B, B)| + 2|E(B^{\triangleright}_A, A)| - |E(A^{\triangleright}_B \cap B^{\triangleright}_A, X)| - |E(Y, A \cap B)|,$$

where $E(Z_1, Z_2)$ denotes, for $Z_1, Z_2 \subseteq V(G)$, the set of edges between Z_1 and Z_2.

Proof. This can be calculated by rearranging sums:

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

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

Lemma 3.3.3. *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:

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

Finally, in order to define tangles of $\vec{S}_k(X)$ and $\vec{S}_k(Y)$ we need to define the notion of *consistency* 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:

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

We define tangles of $\vec{S}_k(Y)$ in G in a similar manner. 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 different notions of 'consistency'.

We are now ready to state the main results of this section. We will show that, with the aid of this order function, we can relate the tangles in $\vec{S}(X)$ to those in $\vec{S}(Y)$.

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

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

Proof of Theorem 5. 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

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

and

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

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 $C_1 \cup C_2 \cup C_3 = Y$.

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 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 $C_1 \cup C_2 \cup C_3 = Y$,

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

Hence, $|E(Z, Y)| < 3k$. Then

$$|(Z, X)|_X = E(Z, Y) < 3k.$$

Thus, $(Z, X) \in \tau$ by (†).

Finally, since $|(A_3, B_3)|_X \leqslant |(C_3, D_3)|_Y < k$ we can conclude by Lemma 3.3.1 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 (†). □

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': that is, 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 $\tau' = {}^{\triangleright}\tau \cap \vec{S}_{4k}(Y)$, and let $\tau'' = {}^{\triangleleft}\tau' \cap \vec{S}_k(X)$. 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 the 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 3.3.4. *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 3.3.3, we have that

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

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

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

In particular, for $x \in A$ we have $|N(x)\cap B^{\triangleright}_A| \leqslant |N(x)\cap A^{\triangleright}_B|$ and thus $x \in (A^{\triangleright}_B)^{\triangleleft}_{B^{\triangleright}_A}$. Similarly, for $x \in B$ we have $|N(x)\cap B^{\triangleright}_A| \geqslant |N(x)\cap A^{\triangleright}_B|$ 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 3.3.5. *For every $x \in A \setminus B$ with $|N(x)\cap B^{\triangleright}_A| > |N(x)\cap A^{\triangleright}_B|$ (equivalently: for every $x \in (B^{\triangleright}_A)^{\triangleleft}_{A^{\triangleright}_B} \setminus B$) we have that $|(A+x, B+x)|_X \leqslant |(A,B)|_X$ and $|(\{x\}, X)|_X \leqslant |(A,B)|_X$.*

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

Proof. We have

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

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

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

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

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

□

Lemma 3.3.6. *If $(A,B)\in\vec{S}(X)$ such that $|(A,B)|_X > |(A,B)^{\triangleright}|_Y$, then their either exists an $x\in(A\setminus B)\cup(B\setminus A)$ 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 an $x\in A\cap B$ such that for $(A',B') = (A-x, B)$ or $(A',B') = (A, B-x)$ we have $|(A',B')|_X < |(A,B)|_X$ and $|(\{x\}, X)|_X \leqslant |(A,B)|_X$.*

Proof. By the proof of Lemma 3.3.4, if $|(A,B)|_X > |(A_B^{\triangleright}, B_A^{\triangleright})|_Y$, then

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

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

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

Every such x is suitable for the x in the assumption by Lemma 3.3.5.

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

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

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

$$2\min\{|N(y) \cap A|, |N(y) \cap B|\} - |N(y) \cap A \cap B| \\ \geqslant 2|N(y) \cap \{x\}| - |N(y) \cap \{x\}| = |N(y) \cap \{x\}|.$$

Hence,

$$|(\{x\}, X)|_X = |N(x)| = \sum_{y \in Y} |N(y) \cap \{x\}| \\ \leqslant \sum_{y \in Y} (2\min\{|N(y) \cap A|, |N(y) \cap B|\} - |N(y) \cap A \cap B|) = |(A, B)|_X.$$

□

We are now ready to prove Theorem 6:

Proof of Theorem 6. Both $\tau'' \cap \vec{S_k}(X)$ 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 $(B, A) \in \tau'' \cap \vec{S_k}(X)$ 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 $|(A, B)|_X = |(A^{\triangleright}_B, B^{\triangleright}_A)|_Y$.

Let $(A', B') = ((A, B)^{\triangleright})^{\triangleleft}$, by Lemma 3.3.4, we have that $A \subseteq A'$ and $B \subseteq B'$. Since $(A', B') \neq (A, B)$ by definition, we can pick $x \in (A' \setminus A) \cup (B' \setminus B)$. Note that $|(A' + x, B' + x)|_X \leqslant |(A, B)|_X$ by Lemma 3.3.5. 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 $|(\{x\}, X)|_X \leqslant |(A, B)|_X$ (again by Lemma 3.3.5) and $(\{x\}, X) \in \tau'' \cap \tau$ this contradicts the fact that $\tau'' \cap \vec{S_k}(X)$, respectively $\tau \cap \vec{S_k}(X)$ are tangles.

If on the other hand $|(A, B)|_X > |(A^{\triangleright}_B, B^{\triangleright}_A)|_X$ then, by Lemma 3.3.6 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 $\tau'' \cap \vec{S_k}(X)$, 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 $(B', A') \in \tau \cap \tau''$ and $\{A, B\}, \{A', B'\}, \{x, X\}$ 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.

3.3.2 Tangles of 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.

We will denote the set of all set separations of the edge set E 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 $\vec{S}(X)$:

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

where $E(v)$ denotes the set of incident edges of v. Again we observe that the degenerate separation (E,E) has maximal possible order and so we shall again implicitly assume that this separation is not contained in any $\vec{S}_k(E)$.

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. In a similar manner 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 $(C,D)^{◀} := (C_D^{◀}, D_C^{◀})$ of X by defining

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

and

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

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

Lemma 3.3.7. *If $(C,D) \leqslant (C',D')$, then $(C,D)^{◀} \leqslant (C',D')^{◀}$*

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

In a similar manner we can define a separation $(C,D)^{▶}$ of Y, however by the symmetry of the situation we will only ever need to talk about the map $(\cdot)^{◀}$.

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 $(C,D)^{◀}$ 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 in $\vec{S}(E)$ 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 $\big((A,B)^E\big)^{\blacktriangleleft} = (A,B)$, 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 3.3.16.

We will show that for a given tangle τ of $\vec{S}_{4k}(X)$, the set $\tau_E \cap \vec{S}_k(E)$ actually is a tangle of $\vec{S}_k(E)$ and dually that, if τ is a tangle of $\vec{S}_{2k}(E)$, then $\tau_X \cap \vec{S}_k(X)$ 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.

As a first step, let us show that the function $|\cdot|_E$ is again submodular. This follows from straightforward calculations, which we nevertheless include here for the sake of completeness:

Lemma 3.3.8. *The order function $|\cdot|_E$ is submodular.*

Proof. As in the proof of Lemma 3.3.1, it is enough to show that, for every $v \in V$, the function

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

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

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

and

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

Summing these two, we get

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

Similarly,

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

and

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

Summing these two, we get

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

Thus, $|\cdot|_v$ is a submodular order 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 3.3.9. *Given a separation (A, B) of X, we have that $|(A, B)|_X \leqslant |(A, B)^E|_E$ and $|(A, B)^E|_E \leqslant 2|(A, B)|_X$.*

Proof. For the first statement we note that:

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

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

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

and thus

$$\begin{aligned}&\sum_{x\in X}\left(2\min\{|E(x)\cap E(A)|,|E(x)\cap E(B)|\}-|E(x)\cap E(A)\cap E(B)|\right)\\=&|E(A)\cap E(B)|.\end{aligned}$$

Since clearly $|(A,B)|_X \geqslant |E(A)\cap E(B)|$, it follows that

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

□

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

Lemma 3.3.10. *Let (C,D) be a separation of E, then $|(C,D)|_E \geqslant |(C^{\blacktriangleleft}_D, D^{\blacktriangleleft}_C)|_X$.*

For the proof of Lemma 3.3.10 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 3.3.10 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 $(C,D)^{\blacktriangleleft}$ if at the end vertex in X of that edge there are fewer incident edges in C than in D.

Lemma 3.3.11. *Let (C,D) be a separation of E and let $e\in E$ be incident with $C^{\blacktriangleleft}_D \smallsetminus D^{\blacktriangleleft}_C$. Then $|(C+e,D-e)|_E \leqslant |(C,D)|_E$ and $(C+e,D-e)^{\blacktriangleleft}=(C,D)^{\blacktriangleleft}$.*

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

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

and

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

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

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 $A \smallsetminus B$ is contained in $C \smallsetminus D$ and every edge incident with $B \smallsetminus A$ is contained in $D \smallsetminus C$:

Lemma 3.3.12. *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 $C^{\blacktriangleleft}_D \smallsetminus D^{\blacktriangleleft}_C$ lies in $C' \smallsetminus D'$ and every edge incident with $D^{\blacktriangleleft}_C \smallsetminus C^{\blacktriangleleft}_D$ lies in $D' \smallsetminus C'$.*

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

Suppose that there exists some such edge e incident with $C^{\blacktriangleleft}_D \smallsetminus D^{\blacktriangleleft}_C$ which does not lie in $C' \smallsetminus D'$ or some such edge e incident with $D^{\blacktriangleleft}_C \smallsetminus C^{\blacktriangleleft}_D$ which does not lie in $D' \smallsetminus C'$. 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 3.3.11 we could then consider the separation $(C' + e, D' - e)$ which must then satisfy $(C' + e, D' - e)^{\blacktriangleleft} = (C', D')^{\blacktriangleleft}$ and $|(C' + e, D' - e)|_E \leqslant |(C', D')|_E \leqslant |(C, D)|_E$, contradicting the choice of (C', D'). □

This observation enables us to perform the necessary calculations to prove Lemma 3.3.10.

Proof of Lemma 3.3.10. By Lemma 3.3.12 we may suppose that every edge incident with $C^{\blacktriangleleft}_D \smallsetminus D^{\blacktriangleleft}_C$ lies in $C \smallsetminus D$ and every edge incident with $D^{\blacktriangleleft}_C \smallsetminus C^{\blacktriangleleft}_D$ lies in $D \smallsetminus C$.

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

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

□

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 3.3.11 allows us to move certain edges from $C \smallsetminus D$ to $D \smallsetminus C$ 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 3.3.7, 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 3.3.13. *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. Let $e = vw$. We observe that, since $v \in C_D^{\blacktriangleleft}$, we have

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

and

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

Thus, $|(C+e, D)|_E \leqslant |(C, D)|_E$. We have $(C, D)^{\blacktriangleleft} \leqslant (C+e, D)^{\blacktriangleleft}$ by Lemma 3.3.7. □

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

Proof. We first note that the set $\tau_X \cap \vec{S}_k(X)$ contains at least one of (A, B) and (B, A) for every separation $(A, B) \in \vec{S}_k(X)$. Indeed, by Proposition 3.3.9 $|(A, B)^E|_E \leqslant 2|(A, B)|_X$, 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 $\{A, B\}$ we have that both (A, B) and (B, A) are contained in $\tau_X \cap \vec{S}_k(X)$. 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 3.3.9, we have that

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

hence $(A, B)^E \in \tau$ or $(B, A)^E \in \tau$. Since $(E(A), E(B))^{\blacktriangleleft} = (A, B)$, 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 $e \in E(B) \smallsetminus E(A)$ since $E(A) = C_1$ and $e \notin C_1$. Thus, $x \in B \smallsetminus A$. Moreover, as $B = {C_2}_{D_2}^{\blacktriangleleft} \subseteq C_D^{\blacktriangleleft}$, we have that $x \in C_D^{\blacktriangleleft}$. Thus, e is incident with $C_D^{\blacktriangleleft}$.

Consequently, applying Lemma 3.3.13 yields $|(C+e, D)|_E \leqslant |(C, D)|_E$. 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), as $(C+e, D)^{\blacktriangleleft} \geqslant (C, D)^{\blacktriangleleft}$ by Lemma 3.3.7 and

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

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

Thus, $\tau_X \cap \vec{S}_k(X)$ is indeed an orientation. That $\tau_X \cap \vec{S}_k(X)$ satisfies the tangle property $(\dagger)$ now follows like this: if $(A_1, B_1), (A_2, B_2), (A_3, B_3)$ would be a triple in $\tau_X \cap \vec{S}_k(X)$ 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 3.3.9. By the above observation τ orients them as $(A_1, B_1)^E$,$(A_2, B_2)^E$ and $(A_3, B_3)^E$, since $\tau_X \cap \vec{S}_k(X)$ 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, A_3. □

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

Theorem 3.3.15. *Given a tangle τ of $\vec{S}_{4k}(X)$, then $\tau_E \cap \vec{S}_k(E)$ is a tangle of $\vec{S}_k(E)$.*

Proof. By Lemma 3.3.10, given some separation $(C, D) \in \vec{S}_k(E)$ we have that $|(C, D)^{\blacktriangleleft}|_X \leqslant |(C, D)|_E$, thus τ contains exactly one of $(C, D)^{\blacktriangleleft}$ and $((C, D)^{\blacktriangleleft})^* = (D, C)^{\blacktriangleleft}$, and consequently $\tau_E \cap \vec{S}_k(E)$ contains exactly one of (C, D) and (D, C), i.e. $\tau_E \cap \vec{S}_k(E)$ is an orientation of $\vec{S}_k(E)$.

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

$$\{(C_1, D_1), (C_2, D_2), (C_3, D_3)\} \subseteq \tau_E \cap \vec{S}_k(E)$$

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

Let $(A_i, B_i) = (C_i, D_i)^{\blacktriangleleft}$ 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 $Z \subseteq B_i = {D_i}_{C_i}^{\blacktriangleleft}$ 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 $C_1 \cup C_2 \cup C_3 = E$,

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

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

$$|(Z, X)|_X = |E(Z, Y)| < 3k.$$

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

Finally, since $|(A_3, B_3)|_X \leqslant |(C_3, D_3)|_E < k$ by Lemma 3.3.10, we can conclude by submodularity that

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

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 (†) 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 (†). □

Putting these two theorems together we immediately obtain that 'double-shifting' a tangle gives, restricted to an appropriate order, a restriction of the original tangle.

Corollary 3.3.16. *Let τ be a tangle of $\vec{S}_{8k}(E)$, then*

$$\tau'' := (\tau_X \cap \vec{S}_{4k}(X))_E \cap \vec{S}_k(E)$$

is a subset of τ.

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

$$\tau''' := (\tau'_E \cap \vec{S}_{2k}(E))_X \cap \vec{S}_k(X)$$

is a subset of τ'.

Proof. By Theorem 3.3.14, 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 $(C, D)^{◀} \in \tau_X \cap \vec{S}_{4k}(X)$ and thus (C, D) is in $(\tau_X \cap \vec{S}_{4k}(X))_E \cap \vec{S}_k(E) = \tau''$. Since τ'' is an orientation of $\vec{S}_k(E)$, we then have that $\tau'' \subseteq \tau$.

For the second part note that, by Theorem 3.3.15, τ''' is a tangle of $\vec{S}_k(X)$. Given $(A, B) \in \vec{S}_k(X) \cap \tau'$ we have, since $((A, B)^E)^{◀} = (A, B)$, that $(A, B)^E$ is in $\tau'_E \cap \vec{S}_{2k}(E)$ and thus $(A, B) \in (\tau'_E \cap \vec{S}_{2k}(E))_X \cap \vec{S}_k(X) = \tau'''$. Since τ''' is an orientation of $\vec{S}_k(X)$, we then have that $\tau''' \subseteq \tau'$. □

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

Corollary 3.3.17 (compare Theorem 5)**.** *Let τ be a tangle of $\vec{S}_{8k}(X)$. Then $\tau' := {}^{▷}\tau \cap \vec{S}_k(Y)$ is a tangle of $\vec{S}_k(Y)$.*

Proof. It is easy to see that $\tau' = (\tau_E \cap \vec{S}_{2k}(E))_Y \cap \vec{S}_k(Y)$ which is a tangle by Theorem 3.3.15 and Theorem 3.3.14. □

Corollary 3.3.18 (compare Theorem 6)**.** *Let τ be a tangle of $\vec{S}_{64k}(X)$, let $\tau' = {}^{▷}\tau \cap \vec{S}_{8k}(Y)$, and let $\tau'' = {}^{◁}\tau' \cap \vec{S}_k(X)$. Then $\tau'' \subseteq \tau$.*

Proof. Consider $\tau_E \cap \vec{S}_{16k}(E)$. By Corollary 3.3.17, we have that

$$\tau' = (\tau_E \cap \vec{S}_{16k}(E))_Y \cap \vec{S}_{8k}(Y).$$

Moreover, again by Corollary 3.3.17 we have that

$$\tau'' = (\tau'_E \cap \vec{S}_{2k}(E))_X \cap \vec{S}_k(X).$$

But now, by Corollary 3.3.16, we note that

$$((\tau_E \cap \vec{S}_{16k}(E))_Y \cap \vec{S}_{8k}(Y))_E \cap \vec{S}_{2k}(E) \subseteq \tau_E \cap \vec{S}_{16k}(E)$$

and thus,

$$(((\tau_E \cap \vec{S}_{16k}(E))_Y \cap \vec{S}_{8k}(Y))_E \cap \vec{S}_{2k}(E))_X \cap \vec{S}_k(X) \subseteq (\tau_E \cap \vec{S}_{16k}(E))_X \cap \vec{S}_k(X),$$

i.e. $\tau'' \subseteq (\tau_E \cap \vec{S}_{16k}(E))_X \cap \vec{S}_k(X)$. Again by Corollary 3.3.16 we have that

$$(\tau_E \cap \vec{S}_{16k}(E))_X \cap \vec{S}_k(X) \subseteq \tau,$$

which shows the claim. □

3.3.3 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.

Regular profiles

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 a *profile* of a system of set separations is a consistent orientation which does not contain any triple of separations of the form

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

and that it is *regular* if it does not contain any co-small separations. It is easy to see that the tangles as defined above are regular profiles, but regular profiles model a broader class of clusters.

Similar statements as in Theorems 5 and 6 can be shown to hold via similar arguments for regular profiles. More precisely, we can show the following with essentially the same proof, which we nevertheless include for the sake of completeness:

Theorem 3.3.19. *Let P be a regular profile of $\vec{S}_{3k}(X)$, then $P' := {}^{\triangleright}P \cap \vec{S}_k(Y)$ is a regular profile of $\vec{S}_k(Y)$.*

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

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

and

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

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

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 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 $i = 1,2,3$. 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 $C_1 \cup C_2 \cup C_3 = Y$,

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

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

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

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

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

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

We also obtain a version of Theorem 6 for regular profiles:

Theorem 3.3.20. *Let P be a regular profile of $\vec{S}_{16k}(X)$, $P' = {}^{\triangleright}P \cap \vec{S}_k(Y)$ and $P'' = {}^{\triangleleft}P' \cap \vec{S}_k(X)$, then $P'' \subseteq P$.*

Proof. Both P'' and $P \cap \vec{S}_k(X)$ are regular profiles of $\vec{S}_k(X)$, suppose that they are distinct. Let $(A,B) \in P$ be a separation of minimal order with the property that $(B,A) \in P'' \cap \vec{S}_k(X)$ 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 $|(A, B)|_X = |(A^{\triangleright}_B, B^{\triangleright}_A)|_Y$.

Let $(A', B') = ((A, B)^{\triangleright})^{\triangleleft}$, by Lemma 3.3.4, we have that $A \subseteq A'$ and $B \subseteq B'$. Since $(A', B') \neq (A, B)$ by definition, we can pick $x \in (A' \setminus A) \cup (B' \setminus B)$. Note that $|(A' + x, B' + x)|_X \leqslant |(A, B)|_X$ by Lemma 3.3.5. 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 $|(\{x\}, X)|_X \leqslant |(A, B)|_X$ (again by Lemma 3.3.5) and $(\{x\}, X) \in P'' \cap P$ this contradicts the fact that $P'' \cap \vec{S_k}(X)$, respectively $P \cap \vec{S_k}(X)$ are profiles.

If on the other hand $|(A, B)|_X > |(A^{\triangleright}_B, B^{\triangleright}_A)|_Y$ then, by Lemma 3.3.6 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 $P'' \cap \vec{S_k}(X)$, respectively $P \cap \vec{S_k}(X)$, is a profile as $(\{x\}, X) \in \tau \cap P''$ and, by the choice of (A, B), either $(A', B') \in \tau \cap P''$ or $(B', A') \in P \cap P''$ and $\{A, B\}, \{A', B'\}, \{\{x\}, X\}$ together contradict the profile property. □

Moreover, we can even show that also the statements about tangles on the edges translate to statements on the profiles of the edges without a lot of changes to the proof:

Proposition 3.3.21. *If P is a regular profile of $\vec{S}_{2k}(E)$, then for $P_X := \{(C, D)^{\blacktriangleleft} : (C, D) \in P\}$, we have that $P_X \cap \vec{S_k}(X)$ is a regular profile of $\vec{S_k}(X)$.*

Proof. If $k \leqslant 1$ we observe that P_X only orients separations (A, B) of order less than 1, thus P_X only orients those separations where every 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 indeed needs to be a profile.

So suppose that $k > 1$.

We first note that the set $P_X \cap \vec{S_k}(X)$ contains at least one of (A, B) and (B, A) for every separation $(A, B) \in \vec{S_k}(X)$. Indeed, by Proposition 3.3.9 we have that $|(A, B)^E|_E \leqslant 2|(A, B)|_X$, 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 that both (A, B) and (B, A) are contained in $P_X \cap \vec{S_k}(X)$. 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 3.3.9, we have that

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

hence $(A, B)^E \in P$ or $(B, A)^E \in P$. Since $(E(A), E(B))^{\blacktriangleleft} = (A, B)$, 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 $D_1 \cap D \subseteq E(A \cap B)$. It is thus easy to see that

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

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

Let x be the end vertex of e in X. Note that $e \in E(B) \smallsetminus E(A)$ since $E(A) = C_1$ and $e \notin C_1$. Thus, $x \in B \smallsetminus A$. Moreover, as $B = C_{2\,D_2}^{\blacktriangleleft} \subseteq C_D^{\blacktriangleleft}$, we have that $x \in C_D^{\blacktriangleleft}$. Thus, e is incident with $C_D^{\blacktriangleleft}$.

Consequently, applying Lemma 3.3.13 yields $|(C+e,D)|_E \leqslant |(C,D)|_E$. 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 $|(\{e\},E)|_E \leqslant 2$ and consequently $(\{e\},E) \in P$, since P is regular and $k > 1$.

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 3.3.7 and

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

since $e \in (D_1 \cap D) \smallsetminus (C_1 \cup C)$.

Thus, $P_X \cap \vec{S}_k(X)$ is indeed an orientation. That $P_X \cap \vec{S}_k(X)$ is a profile now follows like this: if $(A_1,B_1),(A_2,B_2),(B_1 \cap B_2, A_1 \cup A_2)$ would be a triple in $P_X \cap \vec{S}_k(X)$ contradicting the profile property, then P would need to orient $(A_1,B_1)^E$, $(A_2,B_2)^E$ and $(B_1 \cap B_2, A_1 \cup A_2)^E$ by Proposition 3.3.9. 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 $P_X \cap \vec{S}_k(X)$ 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 $P_X \cap \vec{S}_k(X)$ is regular, since if $(X,A) \in P_X \cap \vec{S}_k(X)$, then $(E(X), E(A)) = (E, E(A))$ is a co-small separation in P, which contradicts the regularity of P. □

Proposition 3.3.22. *Given a regular profile P of $\vec{S}_{3k}(X)$, let*

$$P_E := \{(C,D) \in \vec{S}_{3k}(E) \;:\; (C,D)^{\blacktriangleleft} \in P\}.$$

Then $P_E \cap \vec{S}_k(E)$ is a regular profile of $\vec{S}_k(E)$.

Proof. By Lemma 3.3.10, 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 $P_E \cap \vec{S}_k(E)$ contains exactly one of (C,D) and (D,C), i.e. $P_E \cap \vec{S}_k(E)$ is an orientation of $\vec{S}_k(E)$. That the orientation $P_E \cap \vec{S}_k(E)$ is consistent is then immediate from Lemma 3.3.7.

So, it remains to show that $P_E \cap \vec{S}_k(E)$ satisfies the profile property. Let us suppose for a contradiction that there is some set

$$\{(C_1,D_1),(C_2,D_2),(C_3,D_3)\} \subseteq P_E \cap \vec{S}_k(E)$$

such that $(C_1,D_1) \vee (C_2,D_2) = (D_3,C_3)$.

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 $C_1 \cup C_2 \cup C_3 = E$,

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

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

$$|(Z, X)|_X = |E(Z, Y)| < 3k.$$

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 $P_E \cap \vec{S}_k(E)$ is regular, since if $(E, C) \in P_E \cap \vec{S}_k(E)$, then $(E, C)^{\blacktriangleright} = (X, C^{\blacktriangleright}_E)$ is a co-small separation in P, which contradicts the regularity of P. □

It would be nice if one could find a unified result implying both our results about tangles and the results about profiles. Unfortunately, it seems that the nature of the result means that strengthening or weakening the notion of a tangle we consider does not make the statement stronger or weaker, but rather incomparable. Indeed, since we wish to show that tangles in $\vec{S}(X)$ shift to tangles in $\vec{S}(Y)$, 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 it is required to show that a weaker property holds for the shifts, but we only have weaker assumptions on the original orientations.

Forward shifts

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 satisfy (†).

Proposition 3.3.23. *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 $(A,B) \in {}^{\triangleleft}({}^{\triangleright}\tau \cap \vec{S}_{4k}(Y)) \cap \vec{S}_k(X)$ by Theorem 6, and hence $(A,B)^{\triangleleft} = (D,C) \in {}^{\triangleright}\tau \cap \vec{S}_{4k}(Y)$, contradicting the fact that ${}^{\triangleright}\tau \cap \vec{S}_{4k}(Y)$ is a tangle. □

Bipartitions

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 with 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, and 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

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

for bipartitions (A,B) of X and

$$|(C,D)|_Y \coloneqq \sum_{x \in X} \min\{|N(x) \cap C|, |N(x) \cap D|\},$$

for bipartitions (C,D) of Y. These functions are again submodular:

Lemma 3.3.24. *The order functions* $|\cdot|_X$ *and* $|\cdot|_Y$ *are submodular.*

Proof. By symmetry, we only need to show the submodularity of $|\cdot|_X$. Note that it suffices to show that the function

$$|(A,B)|_y \coloneqq \min\{|N(y) \cap A|, |N(y) \cap (B)|\}$$

is submodular for each $y \in Y$, since sums of submodular functions are again submodular.

To that end, let (A,B) and (A',B') be two bipartitions of X and let

$$\begin{aligned} |N(y) \cap A \cap A'| &= p_{aa'}, \quad |N(y) \cap A \cap B'| = p_{ab'} \\ |N(y) \cap B \cap A'| &= p_{ba'}, \quad |N(y) \cap B \cap B'| = p_{bb'}. \end{aligned}$$

so that

$$|(A,B)|_y + |(A',B')|_y = \min\{p_{aa'} + p_{ab'}, p_{ba'} + p_{bb'}\} + \min\{p_{aa'} + p_{ba'}, p_{ab'} + p_{bb'}\}$$

and

$$\begin{aligned} &|(A \cap A', B \cup B')|_y + |(A \cup A', B \cap B')|_y \\ =& \min\{p_{aa'}, p_{ab'} + p_{ba'} + p_{bb'}\} + \min\{p_{aa'} + p_{ba'} + p_{ab'}, p_{bb'}\}. \end{aligned}$$

There are four possible cases for $|(A,B)|_y + |(A',B')|_y$:

$$\begin{aligned} p_{aa'} + p_{ab'} + p_{aa'} + p_{ba'} &= p_{aa'} + (p_{aa'} + p_{ba'} + p_{ab'}); \\ p_{ba'} + p_{bb'} + p_{aa'} + p_{bb'} &\leqslant p_{aa'} + p_{bb'}; \\ p_{aa'} + p_{ab'} + p_{ab'} + p_{bb'} &\leqslant p_{aa'} + p_{bb'}; \\ p_{ba'} + p_{bb'} + p_{ab'} + p_{bb'} &= (p_{ab'} + p_{ba'} + p_{bb'}) + p_{bb'}, \end{aligned}$$

and in each case it is clear that this is at least as big as

$$\begin{aligned} &\min\{p_{aa'}, p_{ab'} + p_{ba'} + p_{bb'}\} + \min\{p_{aa'} + p_{ba'} + p_{ab'}, p_{bb'}\} \\ =&|(A \cap A', B \cup B')|_y + |(A \cup A', B \cap B')|_y. \end{aligned}$$ □

Given some $k \in \mathbb{N}$ we use $\vec{\mathcal{B}}_k(X)$ to denote the set of all bipartitions in $\vec{\mathcal{B}}(X)$ with order lower than k. A *tangle of* $\vec{\mathcal{B}}_k(X)$ then is a consistent orientation τ of $\vec{\mathcal{B}}_k(X)$ which

(BT1) does not contain any separation of the form (A,B) with $|B| \leqslant 1$ or $|N(B)| \leqslant 1$, and

(BT2) where for any three separations $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$ we have that $A_1 \cup A_2 \cup A_3 \neq X$.

Similarly, we define $\vec{\mathcal{B}}_k(Y)$ and tangles of $\vec{\mathcal{B}}_k(Y)$.

We observe that a bipartition (A, B) in $\vec{\mathcal{B}}(X)$ naturally defines a bipartition $(A, B)^{\triangleright} := (C, D)$ in $\vec{\mathcal{B}}(Y)$ by letting

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

and

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

Similarly, a bipartition (C, D) in $\vec{\mathcal{B}}(Y)$ gives rise to a bipartition $(C, D)^{\triangleleft} := (A, B)$ in $\vec{\mathcal{B}}(X)$ via

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

and

$$B := \{x \in X \ : \ |N(x) \cap C| < |N(x) \cap D|\}.$$

Like for set separations, it turns out that, with respect to the order functions defined above, the order of $(A, B)^{\triangleright}$ is at most that of (A, B). In order to show this we will first need the following simple bound on the order of a bipartition.

Lemma 3.3.25. *Let (A, B) be a bipartition of X, then $|(A, B)|_X \leqslant |E(A, Y)|$.*

Proof.

$$|(A, B)|_X = \sum_{y \in Y} \min\{|N(y) \cap A|, |N(y) \cap B|\} \leqslant \sum_{y \in Y} |N(y) \cap A| = |E(A, Y)|. \quad \square$$

This implies that tangles cannot point towards sets of vertices with small neighbourhoods:

Lemma 3.3.26. *If τ is a tangle of $\vec{\mathcal{B}}_k(X)$ and $Z \subseteq X$ with $E(Z, Y) < k$, then $(Z, X \smallsetminus Z) \in \tau$.*

Proof. By Lemma 3.3.25 we have $|(Z', X \smallsetminus Z')|_X \leqslant |E(Z', Y)| \leqslant |E(Z, Y)| < k$ for all $Z' \subseteq Z$, so τ contains an orientation of $(Z', X \smallsetminus Z')$ for all $Z' \subseteq Z$. We show by induction on $|Z'|$ that the orientation in τ is always $(Z', X \smallsetminus Z')$. The induction start for $Z' = \emptyset$ is immediate by (BT1). For the induction step let $Z' \subseteq Z$ with $|Z'| > 0$, pick an arbitrary element $z \in Z'$, and let $Z'' := Z' - z$. By the induction hypothesis, $(Z'', X \smallsetminus Z'')$ is in τ and, by (BT1), $(\{z\}, X - z)$ is in τ. Since $Z'' \cup \{z\} \cup (X \smallsetminus Z') = X$, the tangle τ cannot contain $(X \smallsetminus Z', Z')$ by (BT2) and hence contains $(Z', X \smallsetminus Z')$. $\square$

We can show that the order of the shift of a bipartition is at most the order of the bipartition we started with:

Lemma 3.3.27. *Let $(A, B) \in \vec{\mathcal{B}}(X)$, then $|(A, B)|_X \geqslant |(A, B)^{\triangleright}|_Y$. Similarly, if $(C, D) \in \vec{\mathcal{B}}(Y)$ then $|(C, D)|_Y \geqslant |(C, D)^{\triangleleft}|_X$.*

Proof. Clearly it suffices to prove the first statement by symmetry. Let $(A,B)^{\triangleright} = (C,D)$. Then, by definition, we have $|N(y) \cap A| \geqslant |N(y) \cap B|$ for $y \in C$ and vice versa for $y \in D$. Thus,

$$\begin{aligned}
|(A,B)|_X &= \sum_{y \in Y} \min\{|N(y) \cap A|, |N(y) \cap B|\} \\
&= \sum_{y \in C} |N(y) \cap B| + \sum_{y \in D} |N(y) \cap A| \\
&= |E(C,B)| + |E(D,A)|.
\end{aligned}$$

Conversely, we have

$$\begin{aligned}
|(C,D)|_Y &= \sum_{x \in X} \min\{|N(x) \cap C|, |N(x) \cap D|\} \\
&= \sum_{x \in A} \min\{|N(x) \cap C|, |N(x) \cap D|\} + \sum_{x \in B} \min\{|N(x) \cap C|, |N(x) \cap D|\} \\
&\leqslant \sum_{x \in A} |N(x) \cap D| + \sum_{x \in B} |N(x) \cap C| \\
&= \sum_{x \in A} |E(x,D)| + \sum_{x \in B} |E(x,C)| \\
&= |E(A,D)| + |E(B,C)| \\
&= |(A,B)|_X.
\end{aligned}$$

□

We can now define, given a tangle τ of $\vec{\mathcal{B}}_k(X)$, the shift

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

and, given a tangle τ of $\vec{\mathcal{B}}_k(Y)$, the shift

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

Like for set separations, the shift of a tangle, restricted to a lower order, gives a tangle:

Theorem 3.3.28. *Let τ be a tangle of $\vec{\mathcal{B}}(X)_{4k}$, $\tau' = {}^{\triangleright}\tau \cap \vec{\mathcal{B}}(Y)_k$, then τ' is a tangle of $\vec{\mathcal{B}}(Y)_k$.*

Proof. We first show that τ' is an orientation of $\vec{\mathcal{B}}_k(Y)$. For this, let (C,D) be a bipartition in $\vec{\mathcal{B}}_k(Y)$ and

$$\begin{aligned}
A &= \{x \in X \colon |N(x) \cap C| > |N(x) \cap D|\}, \\
B &= \{x \in X \colon |N(x) \cap C| < |N(x) \cap D|\}, \text{ and} \\
Z &= \{x \in X \colon |N(x) \cap C| = |N(x) \cap D|\}
\end{aligned}$$

then $(C,D)^{\triangleleft} = (A \cup Z, B)$ and $(D,C)^{\triangleleft} = (B \cup Z, A)$.

If both (C,D) and (D,C) are contained in τ', then $(A \cup Z, B)$ and $(B \cup Z, A)$ both are contained in τ, however $A \cap B = \emptyset$, contradicting the fact that τ is a tangle.

If on the other hand neither (C, D) not (D, C) is contained in τ', then neither $(A \cup Z, B)$ nor $(B \cup Z, A)$ is contained in τ and therefore $(B, A \cup Z), (A, B \cup Z) \in \tau$. Moreover, we observe that

$$
\begin{aligned}
|(C, D)|_Y &= \sum_{x \in X} \min\{|N(x) \cap C|, |N(x) \cap D|\} \\
&\geqslant \sum_{x \in Z} \min\{|N(x) \cap C|, |N(x) \cap D|\} \\
&= \sum_{x \in Z} \frac{|N(x) \cap C|}{2} = \frac{|E(Z, Y)|}{2}.
\end{aligned}
$$

Consequently, $|E(Z, Y)| < 2k$. Thus, by Lemma 3.3.26, we have that the separation $(Z, X \setminus Z)$ is contained in τ. However, this contradicts the definition of a tangle, as $A \cup B \cup Z = X$. Thus, τ' is indeed an orientation of $\vec{\mathcal{B}}_k(Y)$.

Secondly, we claim that $(\{y\}, Y - y) \in \tau'$ for every $y \in Y$ with the property that $|(\{y\}, Y - y)|_Y < k$. Indeed, let us suppose for a contradiction that $(Y - y, \{y\}) \in \tau'$. Then, for

$$
A := \{x \in X \colon N(x) = \{y\}\},\ B := X \setminus A,
$$

we have that $(Y - y, \{y\})^{\triangleleft} = (B, A)$. Thus, $(B, A) \in \tau$, however this contradicts the assumption that τ is a tangle, as $|N(A)| = |\{y\}| = 1$.

Additionally, we claim that $(C, D) \in \tau'$ whenever $|N(C)| = 1$ and $|(C, D)|_Y < k$. Again, suppose for a contradiction that $(D, C) \in \tau'$. This implies that the separation $(A, B) := (D, C)^{\triangleright}$ is contained in τ. However, $B \subseteq N(C)$ and thus $|B| \leqslant |N(C)| = 1$ contradicting the fact that τ is a tangle. With that, we have proven that τ' satisfies (BT1).

So, it remains to show that τ' satisfies (BT2). 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 $C_1 \cup C_2 \cup C_3 = Y$.

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 the set $Z = X \setminus (A_1 \cup A_2 \cup A_3)$ is non-empty.

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

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

Hence, $|E(Z, Y)| < 3k$. Thus, by Lemma 3.3.26, we have that $(Z, X \setminus Z) \in \tau$. But, since $|(A_3, B_3)|_X \leqslant |(C_3, D_3)|_Y < k$ and $|E(Z, Y)| \leqslant 3k$ we can conclude that

$$
\begin{aligned}
|(A_3 \cup Z, B_3 \setminus Z)|_X &= \sum_{y \in Y} \min\{|N(y) \cap (A_3 \cup Z)|, |N(y) \cap (B_3 \setminus Z)|\} \\
&\leqslant \sum_{y \in Y} (\min\{|N(y) \cap A_3|, |N(y) \cap B_3|\} + |E(y, Z)|) \\
&< 4k.
\end{aligned}
$$

It follows from (BT1) that $(A_3 \cup Z, B_3 \setminus Z) \in \tau$. However, then

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

and $B_1 \cap B_2 \cap (B_3 \setminus Z) = \emptyset$, contradicting (BT2). □

Like for the tangles of set separations, if we take a tangle on X, shift it to Y and then shift it back to X we obtain a truncation to low order separations of the tangle that we started out with.

Theorem 3.3.29. *Let τ be a tangle of $\vec{\mathcal{B}}(X)_{16k}$, let $\tau' = {}^{\triangleright}\tau \cap \vec{\mathcal{B}}(Y)_{4k}$, and let $\tau'' = {}^{\triangleleft}\tau' \cap \vec{\mathcal{B}}(X)_k$. Then $\tau'' \subseteq \tau$.*

To prove this theorem, we need some lemmas:

Lemma 3.3.30. *If $(A, B) \in \vec{\mathcal{B}}(X)$ and $|(A, B)|_X = |(A, B)^{\triangleright}|_Y$ then $((A, B)^{\triangleright})^{\triangleleft} \geqslant (A, B)$.*

Proof. Let $(C, D) = (A, B)^{\triangleright}$ and $(A', B') = (C, D)^{\triangleleft}$.

By definition, we have that $C = \{y \in Y : |N(y) \cap A| \geqslant \frac{|N(y)|}{2}\}$ and similarly that $D = \{y \in Y : |N(y) \cap A| < \frac{|N(y)|}{2}\}$. The order of (A, B) is thus

$$|E(A, D)| + |E(B, C)|.$$

On the other hand, the order of $(A, B)^{\triangleright}$ is bounded from above by

$$\sum_{x \in A} |E(x, D)| + \sum_{x \in B} |E(x, C)| = |E(A, D)| + |E(B, C)|.$$

Thus, in order for these two orders to be equal we need that every $x \in A$ has at most as many neighbours in D as in C. Since $A' = \{x \in X : |N(x) \cap C| \geqslant \frac{|N(x)|}{2}\}$, this implies that $x \in A'$ for $x \in A$. Thus, $(A, B) \leqslant (A', B') = ((A, B)^{\triangleright})^{\triangleleft}$. □

Lemma 3.3.31. *Let $(A, B) \in \vec{\mathcal{B}}(X)$ and let $(A', B') := ((A, B)^{\triangleright})^{\triangleleft}$, then for every $x \in A' \setminus A$ we have that $|(A + x, B - x)|_X \leqslant |(A, B)|_X$ and moreover $|(\{x\}, X - x)|_X \leqslant 2|(A, B)|_X$.*

Proof. Let $x \in A' \setminus A$ and consider some $y \in Y$. If y is not adjacent to x, it has as many neighbours in A as in $A + x$ and as many neighbours in B as in $B - x$. If y is adjacent to x it has one neighbour more in $A + x$ than in A and one neighbour less in $B - x$ than in B. Thus, by the definition of the order function, we can bound the order of $(A + x, B - x)$ by the order of (A, B) plus the number of vertices in $N(x)$ which have more neighbours in B than in A minus the number of vertices in $N(x)$ which have at least as many neighbours in A as in B, i.e. we have $|(A + x, B - x)|_X \leqslant |(A, B)|_X + |N(x) \cap D| - |N(x) \cap C|$

for $(C, D) := (A, B)^{\triangleright}$. However, as $x \in A'$, we have $|N(x) \cap C| \geqslant |N(x) \cap D|$ and thus $|(A + x, B - x)|_X \leqslant |(A, B)|_X$.

The order of $(\{x\}, X - x)$ is bounded by $|N(x)|$. Observe that the order of (A, B) is at least as big as $|N(x) \cap C|$, as every vertex in C has at least as many neighbours in A as in B. Observe further that $|N(x) \cap C| \geqslant \frac{|N(x)|}{2}$ by the choice of x. Thus, $|(A, B)|_X \geqslant \frac{|N(x)|}{2}$. □

Proof of Theorem 3.3.29. Both τ'' and $\tau \cap \vec{\mathcal{B}}_k(X)$ are tangles of $\vec{\mathcal{B}}_k(X)$, 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 $\leqslant$-maximal. Let $(A', B') = ((A, B)^{\triangleright})^{\triangleleft}$.

If A' is not a subset of A, let $x \in A' \setminus A$ as in Lemma 3.3.31. Then, by the choice of (A, B) and the fact that $(B - x, A + x) \leqslant (B, A) \in \tau''$, we have that $(B - x, A + x) \in \tau$. However, as $|(\{x\}, X - x)|_X \leqslant 2|(A, B)|_X < 2k < 16k$ we have $(\{x\}, X - x) \in \tau$ and $(A, B), (B - x, A + x), (\{x\}, X - x)$ together form a forbidden triple in τ.

Thus, we may suppose that $A' \subseteq A$ and thus $(A', B') \leqslant (A, B)$. Hence, $(A', B') \in \tau$ and therefore $(A, B)^{\triangleright} \in \tau'$ by the definition of τ', as $(A', B') = ((A, B)^{\triangleright})^{\triangleleft} \in \tau$. However, this implies that $(A, B) \in \tau''$, contradicting the fact that $(B, A) \in \tau''$ and τ'' is a tangle. □

Again we note that the proofs of Theorems 3.3.28 and 3.3.29 closely follow those of Theorems 5 and 6, and thus we would suspect that there exists some unified theorem from which both, the theorems about set separations, and those about bipartitions, can be deduced.

However, as with the difference between tangles and profiles, there arises some problems if one tries to do so: in principle, every tangle of set separations induces a tangle of bipartitions. And 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 co-small 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 of each other, although most of the proof strategy is very similar.

3.4 Agile sets

In this section we turn our attention to a different type of dense structures, not directly related to separations or tangles.

Seymour in [70] and, independently, Thomassen in [72] considered, given an ordered set $Z = (x_1, x_2, \dots, x_k, y_1, y_2, \dots, y_k)$ of vertices in a graph G, the question whether there exists a Z-linkage in G. Here, a *Z-linkage* consists of vertex-disjoint paths $P_1, P_2, \dots P_k$ between x_i and y_i. Both, Seymour and Thomassen, independently characterized when a graph contains, for a given ordered set Z of size 4, a Z-linkage.

Now, instead of considering this question for larger k let us generalize the special case of $k = 2$ in a different direction. Let us say that a pair (X_1, X_2) of disjoint vertex sets in G is *independent* if we find two disjoint trees T_1, T_2 in G so that T_1 contains all of X_1 and T_2 contains all of X_2. To the author's knowledge, this notation was first suggested in a mathoverflow post by the anonymous user 'monkeymaths' [66].

Weißauer [76] now used this notion to give a possible definition of when a vertex set is, in some sense, dense in the graph: let us say that a vertex set X is *agile* if for every partition $X = X_1 \dot\cup X_2$ the pair (X_1, X_2) is independent.

A typical example of a large agile set can be found in the complete bipartite graph $K_{2,k}$: the set X of all the degree-2 vertices in this graph is agile, since for any partition $X_1 \dot\cup X_2$ of this set we can obtain disjoint trees T_1 and T_2 by adding one of the degree-k vertices to X_1 and the other to X_2.

Moreover, this notion of agile sets is well-behaved under the minor relation, since if H is a minor of G containing an agile set X, then picking an arbitrary vertex from every branch set corresponding to a vertex in X results in an agile set contained in G.

In light of these observations, Weißauer asked [76] whether, at least qualitatively, a graph contains a large agile set if and only if the graph contains a large $K_{2,k}$ as a minor. More precisely, Weißauer asked the following:

Question 3.4.1. Is it true that for every k there exists an m such that every graph with an agile set of size at least m contains $K_{2,k}$ as a minor?

For $k = 2$ this is the case, since all graphs without a $K_{2,2}$-minor are outerplanar and thus cannot contain an agile set of size 4. This was already observed by Weißauer:

Observation 3.4.2 ([76])**.** If G does not contain a $K_{2,2}$-minor, then G cannot contain an agile set of size $\geqslant 4$, as in that case G is outerplanar.

In the following we will analyse graphs with larger agile sets and answer Question 3.4.1. We will find out that, while the answer to Question 3.4.1 is 'yes' for $k \leqslant 4$, for larger k the question must be answered negatively. However, this is only the case due to one special additional type of graph, and consequently we will be able to show that there is a function $f : \mathbb{N} \to \mathbb{N}$ such that every graph containing an agile set of size $f(k)$ will need to contain a $K_{2,k}$, or this other special type of graph, which is called a regular strip of length k, as a minor.

To achieve this, let us start with a general lemma about graphs containing agile sets, which will allow us to assume that the graph considered is 2-connected:

Lemma 3.4.3. *If G is a graph containing an agile set X of size $l \geqslant 4$, then either G is 2-connected or G contains a proper subgraph still containing an agile set of size l.*

Proof. Suppose G is not 2-connected and let x be a vertex such that the set $\{C_1, \dots, C_k\}$ of components of $G - x$ has size at least 2.

Then there is one component, say C_1, such that all but at most one vertex from X lies in C_1: otherwise pick four vertices v_1, v_2, w_1, w_2 so that neither v_1 and v_2 nor w_1 and w_2 lie in the same component of $G - x$. Then there are no two disjoint trees T_1, T_2 in G such that T_1 contains v_1 and v_2 and T_2 contains w_1 and w_2, as both of these trees would need to contain x, since every path from v_1 to v_2 uses x and also every path from w_1 to w_2 uses x.

Moreover, if v is the only vertex in $X \setminus C_1$, the set $X - v + x$ is again an agile set. It is easy to see that this set is also an agile set of the proper subgraph $C_1 \cup \{x\}$ of G, thus this subgraph contains an agile set of size l. □

While this lemma allows us to essentially assume that every graph containing an agile set is 2-connected, we can, perhaps surprisingly, show something similar for larger connectivity. Of course, not every graph containing a large agile set is 3-connected, as for example $K_{2,k}$ is not, but an agile set still behaves nicely with respect to separations of order 2 or larger. Namely, we can show the following:

Observation 3.4.4. Let G be a graph, (C, D) a separation of G and X an agile set contained in G. Then $X \cap C$ and $X \cap D$ are agile in the corresponding torsos, i.e. in the graphs obtained from $G[C]$ and $G[D]$ by making $C \cap D$ complete.

Proof. Let us show that $X \cap C$ is agile in the torso H corresponding to C. Given any partition $X_1 \dot\cup X_2$ of $X \cap C$, we know that, since X is agile, the pair $(X_1 \cup (X \smallsetminus C), X_2)$ is independent. Thus, we find disjoint trees T_1, T_2 in G so that T_1 contains $X_1 \cup (X \smallsetminus C)$ and T_2 contains A_2. Now clearly $T_1 \cap C$ and $T_2 \cap C$ induce disjoint connected subgraphs of H which contain X_1 and X_2 respectively. Thus, we find the desired trees inside these subgraphs. □

Being able to assume that the graph considered is 2-connected is in fact all we need to answer Question 3.4.1 positively for $k = 3$:

Proposition 3.4.5. *If G does not contain a $K_{2,3}$-minor, then G cannot contain an agile set of size* 5.

Proof. We may assume by Lemma 3.4.3 that G is 2-connected.

Since a graph is outerplanar if and only if the graph contains neither $K_{2,3}$ nor K^4 as a minor, we may suppose that G is either outerplanar or contains K^4 as a minor. Thus, as by Observation 3.4.2 no outerplanar graph contains an agile set of size 5, we may suppose that G contains K^4 as a minor. Thus, (since $\Delta(K^4) = 3$,) G contains K_4 as a topological minor. Since every TK^4 not equal to K^4 contains $K_{2,3}$ as a minor, we therefore may suppose that $K^4 \subseteq G$.

If there is any other vertex $v \in G$, there are two disjoint paths from v to this K^4, since G is 2-connected. However, the K^4 together with v and these 2 paths again include $K_{2,3}$ as a minor, thus there cannot be such a vertex.

Thus, G would need to be equal to K^4, however K^4 does not contain an agile set of size 5. Thus, every graph containing an agile set of size 5 must contain $K_{2,3}$ as a minor. □

Moreover, it is possible to show that Question 3.4.1 is true for $k = 4$:

Proposition 3.4.6. *There exists an m such that, if G does not contain a $K_{2,4}$-minor, then G cannot contain an agile set of size m.*

While Proposition 3.4.6 can be proven directly using either the characterization of graphs without $K_{2,4}$-minor obtained by Ellingham, Marshall, Ozeki, and Tsuchiya in [44], or a result by Dieng ([18]) which states that every graph without a $K_{2,4}$-minor is obtained from an outerplanar graph by the addition of at most 2 vertices, both of these proofs would consist of a rather extensive case distinction. Thus, we will not prove Proposition 3.4.6 directly, instead it will turn up as corollary of a later result.

For $k = 5$ however, Question 3.4.1 needs to be answered negatively, as shown by the following counterexample:

Counterexample 3.4.7. The set of red vertices in the following graph G is agile, however G does not contain a $K_{2,5}$-minor.

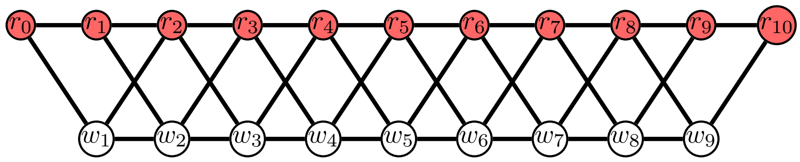

Formally this graph G is constructed as follows: given some $n \in \mathbb{N}$, the vertex set of G consists of the red vertices $r_0, r_1, \dots r_n$ and the white vertices $w_1, \dots w_{n-1}$. The edge set of G consist of an edge between r_i and r_{i+1} for all $0 \leqslant i < n$, an edge between w_i and w_{i+1} for all $1 \leqslant i < n-1$ as well as an edge between any w_i and r_{i+1} and any w_i and r_{i-1} for all $1 \leqslant i < n-1$.

Proof. It is easy to see that the set $X = \{r_0, \dots, r_n\}$ of the red vertices in G is agile.

To see that G does not contain a $K_{2,5}$-minor, suppose for the contrary that G contains a $K_{2,5}$-minor. Then we can find such a minor so that the branch set of every vertex of degree 2 in $K_{2,5}$ consists of only a single vertex of G. Let us denote the branch sets of the vertices of degree 5 of such a $K_{2,5}$-minor in G by H_1 and H_2.

Now consider the set I of those $i \in \{0, \dots, n\}$ for which r_i or w_i corresponds to one of the vertices of degree 2 in $K_{2,5}$. By pigeonhole principle, one of H_1 and H_2 contains, for at least three distinct $i \in I$, neither r_i nor w_i. Let us suppose without loss of generality that H_1 does so and let us denote three such $i \in I$ where H_1 contains neither r_i nor w_i as $i_1 < i_2 < i_3$. Now the set $\{r_{i_2}, w_{i_2}\}$ disconnects every r_j, w_j with $j < i_2$ from every r_k, w_k with $k > i_2$. Therefore, as both, one of r_{i_1}, w_{i_1} and one of r_{i_3}, w_{i_3} correspond to one of the vertices of degree 2 in the $K_{2,5}$, both $\{r_{i_1}, w_{i_1}\}$ and $\{r_{i_3}, w_{i_3}\}$ are adjacent to H_1. But, since H_1 is disjoint from $\{r_{i_2}, w_{i_2}\}$, this contradicts the fact that H_1 is connected, as H_1 would need to meet two components of $G - r_{i_2} - w_{i_2}$. □

We will now prove that this counterexample is essentially the only one:

Theorem 7. *There exists a function $f : \mathbb{N} \to \mathbb{N}$, such that every graph with an agile set of size $f(k)$ either contains $K_{2,k}$ or a regular strip of length k as a minor.*

To prove this, we will heavily rely on a result by Ding which characterizes the graphs not containing $K_{2,k}$ as a minor [32]. Consequently, we need the following definitions by Ding [32]:

We say that a graph G is *internally 3-connected* if we can obtain G from a 3-connected graph by subdividing each edge at most once.

A *fan* is a graph G which consists of a cycle C, three consecutive vertices $a, b, c \in C$ and additional edges between b and some other vertices on C. These additional edges are called the *chords* of the fan. The vertex b is called the *center* of the fan, and the vertices a, b, c are the *corners* of the fan. The *length* of the fan is the number of chords.

Consider graphs G obtained from a cycle C containing two disjoint edges ab and cd by adding some edges between the two distinct paths in $C \setminus \{ab, cd\}$. The added edges are called *chords*. We say that chords f_1f_2 and f_3f_4 *cross* if the four vertices f_1, f_2, f_3, f_4 are pairwise distinct, and they appear in the order f_1, f_3, f_2, f_4, f_1 along C. If given such a graph G where every chord is crossed by at most one other chord, and where if two chords f_1f_2 and f_3f_4 cross, then either f_1f_3 and f_2f_4, or f_1f_4 and f_2f_3 are edges in C, we call G a *strip*. Moreover, we will also call any $H \in \{G - ab, G - cd, G - ab - cd\}$ a *strip* if H has minimum degree at least 2. The *corners* of such a strip are the vertices a, b, c, d and the *length* of the strip is the maximal size of a set of pairwise non-crossing chords with pairwise disjoint endpoints.

Given a graph G, *adding* a fan or strip to G shall mean that we obtain a new graph out of the disjoint union of G and a fan or strip by identifying the corners of the fan or strip with disjoint vertices from G.

Finally, we say that a graph H is an *augmentation* of a graph G if H is obtained from G by adding disjoint fans and strips in such a way that two corners of distinct fans and strips are only allowed to be identified with the same vertex of G if one of them is the center of a fan, and the other one is either a corner of a strip, or also a center of a fan.

We denote, for $m \in \mathbb{N}$, as $\mathcal{A}_m$ the class of all graphs that are augmentations of a graph with at most m vertices, i.e. the class of all those graphs H for which there is a graph G with at most m vertices such that H is an augmentation of G.

A *regular strip* of length k is the graph obtained from two disjoint paths $P_1 := v_1 \dots v_k$, $P_2 := w_1 \dots w_k$ by adding an edge between v_i and w_{i+1} and w_i and v_{i+1} for every $1 \leqslant i < m$. This graph is depicted in the following image:

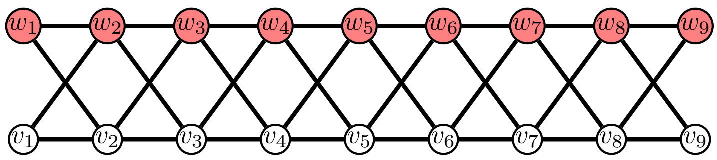

Note that such a regular strip is a strip with corners v_1, w_1, v_k, w_k.

Ding ([32]) now showed the following:

Theorem 3.4.8 ([32, Theorem 5.1], rephrased)**.** *For every $k \in \mathbb{N}$ there is some $m \in \mathbb{N}$ such that every internally 3-connected graph without a $K_{2,k}$-minor is contained in $\mathcal{A}_m$.*

This theorem will allow us to prove Theorem 7. Our proof strategy will be as follows: suppose we are given a graph G which contains a large agile set. We will be able to show that, if G does not contain a $K_{2,k}$-minor, then we find a minor of G which is internally 3-connected and still contains a relatively large agile set, hence we can essentially assume without loss of generality that G is internally 3-connected. We then can assume, using Theorem 3.4.8, that G lies in $\mathcal{A}_m$, thus G is an augmentation of a graph with at most m vertices. Since every such augmentation is obtained by adding a bounded number of fans and strips, this will then imply that one of the fans or strips used for G still contains a relatively large agile set, and we will be able to show that this is only possible for a strip, and that this strip will then contain a regular strip as a minor.

So let us first show that the graphs in $\mathcal{A}_m$ are indeed constructed by adding only a bounded number of fans and strips:

Observation 3.4.9. Every graph in $\mathcal{A}_m$ is obtained from a graph G with at most m vertices by adding at most $\frac{m}{2}$ fans and strips.

Proof. Every vertex of G is a non-center corner of a strip or fan for at most one such fan or strip. Since every strip and fan has at least 2 non-center corners, this gives the desired bound. □

Our next goal is to show that, if an augmentation contains a large agile set, this large agile set cannot be contained in any of the fans used in the construction of this augmentation. Since the corners of such a fan separate the fan from the rest of the graph, this follows from Observation 3.4.4 – as soon as we establish the following:

Observation 3.4.10. Let G be a graph obtained from a fan by making the set of corners complete. Then G cannot contain an agile set of size 7.

Proof. Suppose X is an agile set in G of size at least 7 and let us denote the center of that fan as b. Then there are 4 vertices in X which do not belong to the corners of the fan. Let us denote them as v_1, v_2, v_3, v_4 and assume that they lie in this order on the cycle C used in the construction of G.

Since X is agile, there are disjoint paths P_1 from v_1 to v_3 and P_2 from v_2 to v_4. However, as $\{b, v_2, v_4\}$ together separate v_1 from v_3 in G and P_1 can neither contain v_2 nor v_4, it needs to be the case that b is contained in P_1. But similarly, $\{b, v_1, v_3\}$ separates v_2 from v_4 in G and thus b is contained in P_2, contradicting the fact that P_1 and P_2 are disjoint. □

Corollary 3.4.11. *Let G be a graph containing a fan H as a subgraph such that the corners of H separate the rest of H from the rest of G. If G contains an agile set X, then X cannot contain more than* 6 *vertices from H.*

Proof. Immediate from Observation 3.4.4 and Observation 3.4.10. □

Next we would like to show, given an augmentation G and a strip used in its augmentation process, that, if the vertices of that strip in G contain a large agile set, then this strip needs to contain a large regular strip as a minor.

For this we first observe that a strip contains a regular strip as a minor, as soon as the strip has enough pairs of crossing chords:

Lemma 3.4.12. *Let G be a strip containing k distinct pairs of crossing chords. Then G contains a regular strip of length k as a minor.*

Proof. Let C be the cycle used in the construction of G, and let ab and cd be the two edges of C for which we added chords between $C \smallsetminus \{ab, cd\}$. Let us denote as P_1 and P_2 the two paths together forming $C \smallsetminus \{ab, cd\}$, where P_1 starts in a and P_2 starts in b.

Let us denote the pairs of crossing chords as $\{v_1^i w_1^i, v_2^i w_2^i\}$, where v_j^i is contained in P_1 and we enumerate these pairs so that v_1^i appears before both, v_2^i and v_1^{i-1} on P_1, for every $1 \leqslant i < k$. Then, since every chord in a strip crosses at most one other chord, we have that v_2^i appears before v_1^{i+1} on P_1 (or is identical to v_1^{i+1}) and that the relation between the w_j^i on P_2 is such that w_2^i appears before w_1^i which appears before w_2^{i+1} on P_2 (again, or that $w_1^i = w_2^{i+1}$), for every $1 \leqslant i < k$.

Thus, suppressing every vertex on P_1 or P_2 which is not one of the v_j^i or w_j^i and then contracting every existing edge between v_2^i and v_1^{i+1} as well as every edge between w_1^i and w_2^{i+1} for every $1 \leqslant i < k$ gives the desired regular strip of length k. □

With this lemma at hand we can now show the following:

Lemma 3.4.13. *There is a function $j : \mathbb{N} \to \mathbb{N}$, namely $j(k) := 22(k-1)+4k+1$, such that whenever the graph G obtained from a strip H by making the set of corners of H complete contains an agile set of size $j(k)$, then H contains a regular strip of length k as a minor.*

Proof. By Lemma 3.4.12, if H contains at least k pairs of crossing chords, then H contains a regular strip of length k as a minor. So suppose that H contains at most $k-1$ pairs of crossing chords. Let X be the set of vertices incident with the edges of these chords together with the four corners of the strip. It is easy to see that $G - X$ contains at most $2(k-1)$ components, and that each of these components is either a path or a strip without crossing chords. Moreover, each of these components is adjacent to at most 4 vertices in X. By the pigeonhole principle, one of these components needs to contain, since $j(k) > 22(k-1) + 4k$, at least 11 vertices of our agile set. Let Y be the vertex set of one such component, let K' be the subgraph of G induced on $Y \cup N(Y)$ and denote as K the graph obtained from K' by adding all edges between the vertices in $N(Y)$. Since $N(Y)$ separates Y from the rest of G, it is, by Observation 3.4.4, enough to show that K does not contain an agile set of size at least 11.

The *ladder of length n* is the graph on the set $[n] \times \{0, 1\}$ where we add an edge between (x, y) and (x', y') precisely if $|x - x'| + |y - y'| = 1$. The *endvertices* of a ladder of length n are the four vertices $(0, 0), (0, 1), (n-1, 0)$ and $(n-1, 1)$.

We claim that K is a minor of a graph obtained from a ladder of large enough length by making the 4 endvertices of this ladder complete, i.e. we claim that K is a minor of the type of graph depicted in Fig. 3.1. Indeed, if $G[Y]$ is a strip without crossing chords, then the graph obtained from K' by removing the edges between the vertices from $N(Y)$ is also such a strip. Moreover, the corners of this strip are the vertices from $N(Y)$. Now we find this strip as a minor in a large enough ladder, with the additional property that the branch sets of the vertices from $N(Y)$ each contain one of the endvertices of the ladder. Consequently, K is a minor of the graph obtained from that ladder by adding all edges between the endvertices of that ladder.

If on the other hand $G(Y)$ is a path, then again we find the graph obtained from K' by removing the edges between the vertices from $N(Y)$ as a minor in a

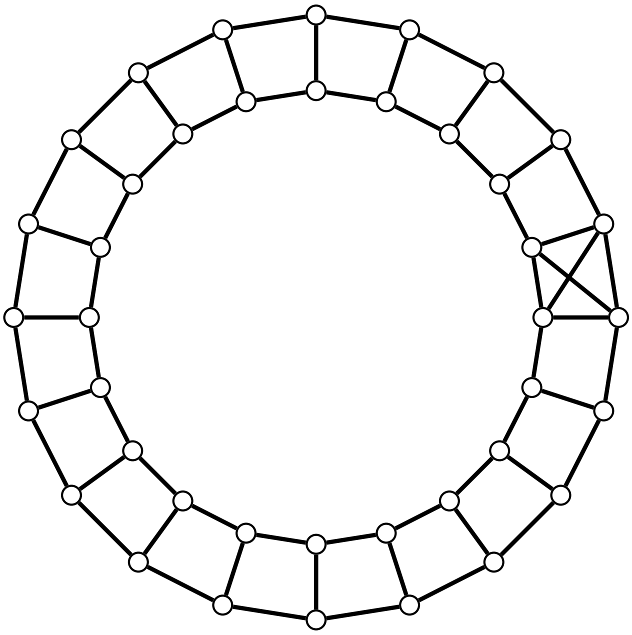

Figure 3.1: The type of graph of which K is a minor.

large enough ladder, with the additional property that the branch sets of $N(Y)$ each contain an endvertex from that ladder.

But, for any $n \in \mathbb{N}$, the graph obtained from a ladder of length n by making its endvertices complete cannot contain an agile set of size 11: by pigeonhole principle, one of the rails (that is one of the sets $\{(i,0) : i \in [n]\}, \{(i,1) : i \in [n]\}$) needs to contain 6 vertices from our agile set. But if we partition these vertices alternatingly, we see that is not possible to connect the two partition classes disjointly.

Since containing an agile set is a minor-closed property, K can thus also not contain an agile set of size 11 contradicting, by Observation 3.4.4, the assumption that Y contains 11 vertices from our agile set. □

Corollary 3.4.14. *Let G be a graph containing a strip H as a subgraph, such that the corners of H separate the rest of H from the rest of G. If G contains an agile set X containing more than $j(k)$ many vertices from H, then G contains a regular strip of length k as a minor.*

Proof. Immediate from Observation 3.4.4 and Lemma 3.4.13. □

We are now ready to use Theorem 3.4.8 to show that every internally 3-connected graph containing a large agile set will indeed need to contain a large $K_{2,k}$ or a large regular strip as a minor:

Lemma 3.4.15. *There exists a function $g : \mathbb{N} \to \mathbb{N}$ such that every internally 3-connected graph which contains an agile set of size $g(k)$ contains $K_{2,k}$ or a regular strip of length k as a minor.*

Proof. By Theorem 3.4.8, we find some $m \in \mathbb{N}$ such that every graph without a $K_{2,k}$-minor is contained in $\mathcal{A}_m$. Let $g(k) := m + \frac{m}{2}\max\{j(k), 6\}$. Now let G be a graph containing an agile set of size $g(k)$. By Theorem 3.4.8 we know that G either contains a $K_{2,k}$ as a minor or is contained in $\mathcal{A}_m$. By Observation 3.4.9, in the second case G is obtained from a graph of at most m vertices by augmenting at most $\frac{m}{2}$ fans and strips. By the pigeonhole principle, one of the augmented fans or strips needs to contain an agile set of size at least $\max\{j(k), 6\}$. By Corollary 3.4.11, this cannot be a fan, so it needs to be a strip. However, by Corollary 3.4.14, this implies that G contains a regular strip of length k as a minor. □

In order to extend this result to graphs that are not locally 3-connected, it will be essential to analyse how separations of order 2 in a graph containing a large agile set can behave. As it turns out, we are able to assume that they are all pairwise nested:

Lemma 3.4.16. *Let $n \in \mathbb{N}$ and let G be a graph which is minor-minimal with the property of containing an agile set X of size n, then the separations of order 2 in G form a nested set.*

Proof. Suppose that two separations (A, B) and (C, D) of order 2 in G cross. If X meets all quadrants $V \setminus (B \cup D), V \setminus (A \cup D), V \setminus (B \cup C)$ and $V \setminus (A \cup C)$, then this would contradict the agility of X by partitioning vertices in opposite quadrants into the same partition class.

Thus, at least one quadrant, say without loss of generality $V \setminus (B \cup D)$, contains no vertex from X. Then, by minor-minimality of G, the quadrant contains no vertex, since contracting any edge adjacent to such a vertex results in a minor of G in which X is still agile. Thus, we may assume that $G[A \cap C]$ consists of just an edge between the sole vertices $v \in (A \cap B) \setminus D$ and $w \in (C \cap D) \setminus B$.

We denote the second vertex in $A \cap B$ as v' and the second vertex in $C \cap D$ as w'.

We claim that we can contract the edge vw, contradicting the minimality of G.

Suppose first, that at most one of v, w is contained in X, say $v \notin X$. Suppose that X is not agile in $G' = G/vw$ and let us denote the partition of X which witnesses this as $X = X_1 \dot\cup X_2$. Since X is agile in G, there are connected disjoint subgraphs G_1, G_2 of G such that $X_1 \subseteq G_1$ and $X_2 \subseteq G_2$. Moreover, we may assume without loss of generality that $V(G_1) \cup V(G_2) = V$ and that $v \in G_1$ and $w \in G_2$, as otherwise $G'[V(G_1)]$ and $G'[V(G_2)]$ are also connected disjoint subgraphs of G'.

Now if $w' \in G_1$, then $B \cap C \supseteq G_2$. Then, for $V_1' = V(G_1) - v$ and $V_2' = V(G_2) + v$, we have that $G'[V_1']$ and $G'[V_2']$ are connected and contain X_1 and X_2, respectively. Thus, $X_1 \dot\cup X_2$ does not witness that X is not agile in G'.

So suppose $w' \in G_2$. Then $v' \in G_2$, since G_2 is connected and $\{v, v'\}$ separates w from w'. By a symmetric argument to the above we may assume that $w \in X$, as otherwise $G'[V(G_1) + w]$ and $G'[V(G_2) - w]$ are connected and contain X_1 and X_2, respectively. Thus, we have that $X_1 \subseteq V \setminus (D \cup A)$.

If $X \cap B \cap C \subseteq X_1$, the partition $X = X_1 \dot\cup X_2$ would again not witness that X is not agile in G', as both $G'[B \cap C]$ and $G'[V \setminus (B \cap C)]$ are connected. So we may suppose that $X_2 \cap (B \cap C)$ is non-empty.

Moreover, there do not exist connected subgraphs G_1'' and G_2'' of $G[B \cap C]$ such that $v \notin G_1''$ and such that $X_1 \subseteq G_1''$ and $X_2 \cap (B \cap C) \subseteq G_2''$, as we could otherwise replace $G_1 \cap B \cap C$ and $G_2 \cap B \cap C$ with these subgraphs which then shows that $X_1 \dot\cup X_2$ is not a partition witnessing that X is not agile in G'.

Additionally, X is not completely contained in C, as G was chosen $\subseteq$-minimal and $X \cap C$ is agile in the torso obtained from $G[C]$ by Observation 3.4.4, and this torso is a proper minor of G since there exists the vertex $v' \in (A \cap B) \setminus C$.

Thus, let $x \in X \setminus C$ and let $X_1' = X_1 + x$ and $X_2' = X_2 - x$. Since X is agile in G, there are disjoint connected subgraphs T_1' and T_2' of G with $X_1' \subseteq T_1'$ and $X_2' \subseteq T_2'$. Now $w' \in T_1'$, since $\{w, w'\}$ separates X_1' in G and $w \in X_2$. On the other hand $v \in T_1'$ since X_1', X_2' look the same on $G[B \cap C]$ as X_1, X_2. But then T_2 cannot connect w to $X_2 \cap (B \cap C)$, which is a contradiction.

It thus remains the case that $v, w \in X$. Then $v', w' \notin X$, as $\{v, v'\}$ separates w from $V \setminus A$ and $\{w, w'\}$ separates v from $V \setminus C$ and both, $V \setminus A$ and $V \setminus C$ need to contain vertices from X by the minor minimality of G.

We may assume that X meets both $V \setminus (A \cup D)$ and $V \setminus (C \cup B)$ in vertices x and y, say, as otherwise the corresponding quadrant would contain only an edge between v and w', respectively w and v', which could be contracted by the previous argument.

Now consider a partition $X_1' \dot\cup X_2'$ of X where $w, x \in X_1'$ and $v, y \in X_2'$. This partition witnesses that X is not agile, which is a contradiction. □

Thus, the set of all separations of order 2 of such a minor-minimal G which are neither small nor co-small form a tree-set, which in turn gives us a tree-decomposition $(T, \mathcal{V})$ of G along all these separations. In particular all torsos of this tree-decomposition are 3-connected. If a large subset of our agile set is contained in one of the torsos of this decomposition, we know by Observation 3.4.4, that it is still agile in the torso, and since the torso is 3-connected, we can then apply Lemma 3.4.15 to deduce that the torso, and thus also the original graph, contains a large $K_{2,k}$ or a large regular strip as a minor.

However, this does not need to be the case. But we can use the structure given by T to analyse the case where it fails. Namely, if no torso contains a large agile set, then our agile set needs to be spread across a lot of different bags of the decomposition. This can either be the case in a 'path-like' or in a 'star-like' way. Let us first deal with the 'path-like' case:

Lemma 3.4.17. *There exists a function $h : \mathbb{N} \to \mathbb{N}$ such that the following holds: if G is a graph containing an agile set X and a sequence*

$$(A_1, B_1) \leqslant \ldots \leqslant (A_{h(k)}, B_{h(k)})$$

of separations of order 2, so that $(B_i \setminus B_{i+1}) \cap X \neq \emptyset$ for all $1 \leqslant i < h(k)$, then G contains a $K_{2,k}$ or a regular strip of length k as a minor.

Proof. Let $x_i \in B_i \setminus B_{i+1} \cap X$ and note that the x_i are pairwise distinct. Let $A_i \cap B_i = \{s_i^1, s_i^2\}$.

If we consider the partition of X given by the two classes $X_1 = \{x_1, x_3, \ldots\}$ and $X_2 = \{x_2, x_4, \ldots\}$, and corresponding disjoint trees T_1 and T_2 containing X_1 and X_2, respectively, we observe that there need to be, for every $2 < i < n - 3$, two disjoint paths from $A_i \cap B_i$ to $A_{i+1} \cap B_{i+1}$ in $B_i \cap A_{i+1}$, one contained in T_1 and the other contained in T_2.

We say that the pair $\{i, (i+1)\}$ is *free* if there are two pairs of such paths T_1, T_2, where one consists of a path between s_i^1 and s_{i+1}^1 and a path between s_i^2 and s_{i+1}^2, and the other consists of a path between s_i^1 and s_{i+1}^2 and a path between s_i^2 and s_{i+1}^1. Otherwise, the pair $i(i+1)$ is said to be *restrictive*.

We note that, if there are, for $2 < i < n-2$, two consecutive pairs $\{(i-1), i\}$ and $\{i, (i+1)\}$ which both are restrictive, then for one of the pairs, say $\{i, (i+1)\}$, there need to be two pairs $T_{1,i}, T_{2,i}$ and $T'_{1,i}, T'_{2,i}$ of disjoint trees in $B_i \cap A_{i+1}$, such that $T_{1,i}$ contains s_i^1, x_i, s_{i+1}^1 and $T_{2,i}$ contains s_i^2, s_{i+1}^2 and $T'_{1,i}$ contains s_i^1, s_{i+1}^1 and $T'_{2,i}$ contains s_i^2, x_i, s_{i+1}^2. This is due to the fact that X is agile, and we can thus consider the partition of X obtained from the one above by changing only the class to which x_i belongs. Moreover, we may suppose without loss of generality that $T_{2,i}$ and $T'_{1,i}$ are paths.

If there are, for the pair $\{i, (i+1)\}$, these four trees $T_{1,i}, T_{2,i}, T'_{1,i}, T'_{2,i}$ as above with the additional property that the two paths $T_{2,i}$ and $T'_{1,i}$ meet, then we say that the restrictive pair $\{i, (i+1)\}$ is *weakly free*.

We may suppose that $h(k)$ is chosen such that there either is a large interval $l < i < m$ with the property that every pair $\{i, (i+1)\}$ from that interval is restrictive and not weakly free, or that there is a large collection of pairs which are all free or weakly free.

In the former case, by the definition of weakly free and the above observation about two adjacent restrictive pairs, we may suppose that m (and thus $h(k)$) is chosen such that there are at least k pairs $\{i, (i+1)\}$ from the interval $l < i < m$, for which we find trees $T_{1,i}, T_{2,i}, T'_{1,i}, T'_{2,i}$ as above with the additional property that $T_{2,i}$ and $T'_{1,i}$ are disjoint. In that case we find a $K_{2,k}$-minor in G as follows: the branch sets for the two vertices of degree k each consist of a path between $A_l \cap B_l$ and $A_m \cap B_m$ formed by concatenating the $T_{2,i}$'s and $T'_{1,i}$'s, respectively. Now we find a path between x_i and $T_{2,i}$ and a path between x_i and $T'_{1,i}$ both contained in $T_{1,i} \cup T'_{2,i}$ and thus both contained in $B_i \cap A_{i+1}$. The union of these two paths form, for each of the k restrictive pairs considered, the branch set of a vertex of degree 2 in our $K_{2,k}$-minor.

So we may suppose that there is a collection of at least k pairs $\{i, (i+1)\}$ which are all free or weakly free.

In this case we claim that whenever a pair $\{i, (i+1)\}$ is free or weakly free, there exists a *cross* in $B_i \cap A_{i+1}$. Here, a *cross* shall be a graph consisting of two disjoint paths P_1 and P_2 between s_i^1 and s_{i+1}^1 and s_i^2 and s_{i+1}^2 together with two disjoint paths Q_1 from P_1 to P_2 and Q_2 from P_2 to P_1, such that Q_1 starts in a vertex before the endvertex of Q_2 and Q_2 starts in a vertex before the endvertex of Q_1.

To see that such a cross exists if the pair $\{i, (i+1)\}$ is free we observe the following: since $\{i, (i+1)\}$ is free, there are two disjoint paths $P_{1,i}, P_{2,i}$, where $P_{1,i}$ is a path between s_i^1 and s_{i+1}^1 and $P_{2,i}$ is a path between s_i^2 and s_{i+1}^2. Additionally, there are two disjoint paths $P'_{1,i}, P'_{2,i}$ where $P'_{1,i}$ is a path between s_i^1 and s_{i+1}^2 and $P'_{2,i}$ is a path between s_i^2 and s_{i+1}^1. Let us suppose that these four paths are chosen so that their union contains as few edges as possible. We set $P_1 = P_{1,i}$ and $P_2 = P_{2,i}$. Now both $P'_{1,i}$ and $P'_{2,i}$ contain subpaths between P_1 and P_2 which are, except of their endvertices, disjoint from P_1 and P_2 and it follows from the choice of $P_{1,i}, P_{2,i}, P'_{1,i}, P'_{2,i}$ that some two of these subpaths need to form our desired paths Q_1 and Q_2.

If on the other hand the pair $\{i, (i+1)\}$ is weakly free, there are two pairs

$T_{1,i}, T_{2,i}$ and $T'_{1,i}, T'_{2,i}$ of disjoint trees in $B_i \cap A_{i+1}$, such that $T_{1,i}$ contains s_i^1, x_i, s_{i+1}^1 and $T_{2,i}$ contains s_i^2, s_{i+1}^2 and $T'_{1,i}$ contains s_i^1, s_{i+1}^1 and $T'_{2,i}$ contains s_i^2, x_i, s_{i+1}^2. Moreover, $T_{2,i}$ and $T'_{1,i}$ are paths which meet in a common vertex y. Again, by choosing these trees so that there union is as small as possible, we observe that we indeed need to find the desired cross.

Thus, each such pair give rise to a cross, and it is easy to see that if we combine all these crosses we obtain a regular strip of length k as a minor in G. □

We are now ready to prove Theorem 7:

Theorem 7. *There exists a function $f : \mathbb{N} \to \mathbb{N}$, such that every graph with an agile set of size $f(k)$ either contains $K_{2,k}$ or a regular strip of length k as a minor.*

Proof of Theorem 7. By Lemma 3.4.3 we may suppose that G is 2-connected. Moreover, we may suppose that G is minor-minimal with the property that G contains an agile set X of size $f(k)$.

By Lemma 3.4.16, the regular separations of order 2 in G form a nested set, thus by Lemma 2.6.2, there is an S-tree T corresponding to this set. This S-tree corresponds to a tree-decomposition, and we consider the parts of this decomposition containing vertices of our agile set. If there is one node of T whose corresponding part contains at least $g(k)$ many vertices of our agile set, then, since the torso of this part is 3 connected, this torso and thus G contains a $K_{2,k}$ or a regular strip of length k as a minor by Lemma 3.4.15.

So each part of this decomposition contains less than $g(k)$ many vertices of our agile set. Thus, by the pigeonhole principle, we can choose $f(K)$ such that we either find a sequence of separations of G as in Lemma 3.4.17, or that there is a node t if T such that there are $n >> k$ many different components of $T - t$ which each contain a vertex t' such that the part of the tree-decomposition corresponding to t' contains a vertex of our agile set which is not contained in the part corresponding to t.

In the first case, we are immediately done by Lemma 3.4.17, so suppose that there indeed is a node t of T such that there are $n >> k$ different components of $T - t$ which each contain a vertex t' such that the part of the tree-decomposition corresponding to t' contains a vertex of our agile set which is not contained in the part corresponding to t.

Let us denote the separations corresponding to the incoming edges from these components to t as $(A_1, B_1), \dots, (A_n, B_n)$ and note that they form a star.

We now ask whether $A_i \setminus B_i$ contains at least two vertices of X, or if this set contains only one such vertex. If at least k of the sets $A_i \setminus B_i$ contain at least two vertices from our agile set, we consider a partition $X_1 \dot\cup X_2$ of X where, for each of these $A_i \setminus B_i$, we add one vertex from $X \cap A_i \setminus B_i$ to X_1 and all others to X_2. This results in two trees T_1 and T_2, where each of the separators $A_i \cap B_i$ needs to contain one vertex from T_1 and one vertex from T_2. Moreover, T_1 and T_2 still need to be connected after deleting all the sets $A_i \setminus B_i$ which contain two vertices from X. These two connected sets form the two vertices of high degree of a $K_{2,k}$. The vertices of degree 2 can then by obtained from the components of the sets $A_i \setminus B_i$, since each such component needs, as G is 2-connected, to send an edge to both vertices in $A_i \cap B_i$ and thus sends an edge to both T_1 and T_2.

So suppose that at least $g(k)$ of the sets $A_i \setminus B_i$ only contain one vertex from X. Then we know, since G was chosen minor-minimal, that $A_i \setminus B_i$ consists for each such i of just this one vertex from X and that this vertex is, as G is 2-connected, adjacent to both vertices in $A_i \cap B_i$. Note that no two separations corresponding to incoming edges of t can have the same separator, since then the supremum of these two separations would also lie in our nested set, and would imply that t has only degree 3 in T.

Let us show that, if we contract, for every incoming edge to t which corresponds to a separation (A, B), one component of the set $A \setminus B$ down to a single vertex, and delete all other components of $A \setminus B$, we are left with an internally 3-connected graph. For this we only need to show that G does not contain any edge e in $C \cap D$, as the torso corresponding to t is 3-connected. So suppose G does contain an edge $e \in C \cap D$. Then, since G was chosen minor-minimal, deleting this edge results in X no longer being agile, say because of the partition $X_1 \dot\cup X_2$, which was independent in G as witnessed by T_1 and T_2, and suppose that T_1 contains e. Then T_2 would need to be contained entirely in $A \setminus B$, as otherwise T_2 is disjoint from $A \setminus B$, and thus replacing e with a path between the two vertices in $A \cap B$ contained in A results in (X_1, X_2) being independent in $G - e$. Thus, since T_2 is contained in $A \setminus B$, we may assume that T_2 consists of just one vertex x from X, and $A \setminus B = \{x\}$. But now, since (X_1, X_2) is not independent in $G - e$, this implies that x is a separator in $G - e$. But the only neighbours of x are the vertices in $A \cap B$, and thus, since G contains more than 3 vertices, one of the two vertices in $A \cap B$ would be a separator of G, contradicting Lemma 3.4.3.

Thus, if we contract for every incoming edge to t with separation (A, B) one component of the set $A \setminus B$ down to a single vertex, and delete all other components of $A \setminus B$, we are left with an internally 3-connected graph. This graph still needs to contain an agile set of size $g(k)$, since the set of all the vertices from X which are the unique vertex from X in one of the sets $A_i \setminus B_i$ is agile in this restricted graph. Thus, by Lemma 3.4.15, we again find a $K_{2,k}$ or a regular strip of length k as a minor. □

Now Theorem 7 also gives a proof of Proposition 3.4.6:

Proof of Proposition 3.4.6. If G contains a large enough agile set, then G contains, by Theorem 7, either a $K_{2,4}$ or a regular strip of length 4 as a minor. Such a strip however also contains $K_{2,4}$ as a minor, which proves Proposition 3.4.6. □

3.4.1 Generalizations

Let us now look at possible variations of the notion of an agile set. One such natural variation is the following: instead of just partitioning our set X into two subsets, we might allow partitions into more partition classes and try to connect the vertices in each of these classes disjointly. More precisely, let us say that, given a graph $G = (V, E)$ and an integer m, a set $X \subseteq V$ is *m-agile* in G if for every partition $X = X_1 \dot\cup \ldots \dot\cup X_m$ (where we allow empty partition classes) there are vertex-disjoint connected subgraphs $T_1, \ldots, T_m \subseteq G$ such that $X_i \subseteq T_i$. So, X is 2-agile if and only if X is agile. If a set X is m-agile for every m, we say that X is *dexterous*. Note that this is equivalent to being $\left\lceil \frac{|A|}{2} \right\rceil$-agile.

Again, containing an m-agile or a dexterous set is closed under the minor relation in that, if H is a minor of G and H contains an m-agile or dexterous set of size k, say, then G also contains an m-agile or dexterous set of size k.

We again try to characterize, qualitatively, the existence of a large m-agile or dexterous set via a minor. A natural graph containing an m-agile set of size k is the complete bipartite graph $K_{m,k}$. On the other hand the complete graph K^m contains a dexterous set of size m.

Another example of such graphs can be found in grids and, since neither K^5 not $K_{3,3}$ is a minor of the grid, these grids are another class of graphs containing a large dexterous or m-agile set. For dexterous sets we need to take a quadratic grid:

Example 3.4.18. In the $N^2 \times N^2$–grid taking every N^{th} vertex of the diagonal gives a dexterous set of site N.

Of course, since every dexterous set is m-agile for every m, the quadratic $N^2 \times N^2$-grid also contains an m-agile set of size N. But such a set can actually already be found in a rectangular grid where the small side just needs to have size $2m - 1$:

Proposition 3.4.19. *The* $(2m - 1) \times ((N - 1)m + 1)$ *grid contains an* m*-agile set of size* N*.*

Proof. Let us denote the vertices of the grid by

$$\{v_{i,j} \;:\; 1 \leqslant i \leqslant 2m - 1, 1 \leqslant j \leqslant N(m - 1)\}.$$

We claim that the set $\{v_{m-1,jk+1} \;:\; 0 \leqslant j \leqslant N - 1\}$ is m-agile. How to construct the required trees is illustrated in the following picture:

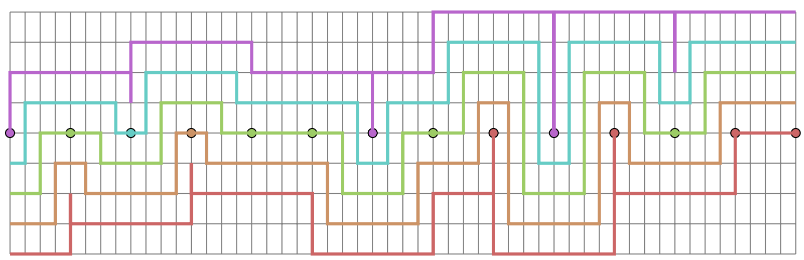

□

In particular, as every m-agile set of size $2m$ is also dexterous, also the $(N - 1) \times (\frac{(N-1)N}{2} + 1)$-grid contains a dexterous set of size N.

If we only seek for a quantitative result, not only in terms of the size of our agile set but also in terms of the m for which our set is m-agile, we can actually show that conversely a large enough l-agile set, for large enough l, forces the existence of either a $K_{m,N}$ or a rectangular grid as a minor. Similarly, we can show that a large enough dexterous set forces the existence of a large complete graph or a large quadratic grid as a minor.

For dexterous sets this can be obtained immediately, as the existence of a large dexterous set implies that our graph has high tree-width. This can either be shown directly or by using a result by Diestel, Jensen, Gorbunov and Thomassen ([28]) about m-connected sets. Following their definition, a vertex set X is *m-connected* if $|X| \geqslant m$ and for any two subsets $X_1, X_2 \subseteq X$ with $|X_1| = |X_2| \leqslant m$ one can find $|X_1|$ many disjoint paths between X_1 and X_2.

This notion of m-connected vertex sets is related to our m-agile sets in that an m-agile set needs to be m-connected:

Proposition 3.4.20. *If a graph contains an m-agile set Z of size at least m, then Z is m-connected.*

Proof. Let $X, Y \subseteq Z$ such that $|X| = |Y| \leqslant m$. Then we can find pairs $\{x_1, y_1\}, \dots \{x_{|X|}, y_{|X|}\}$ so that $x_i \in X$ and $y_i \in Y$ and so that $x_i = y_j$ only if $i = j$. We now construct a partition of Z into classes $X_1, \dots, X_{|X|}$ by defining $X_i = \{x_i, y_i\}$ whenever $i \geqslant 2$ and $X_1 = X \setminus \{x_2, y_2, \dots, x_{|X|}, y_{|X|}\}$.

Since Z is m-agile, we thus find disjoint trees $T_1, \dots, T_{|X|}$ so that $X_i \subseteq T_i$. In particular, T_i contains a path between x_i and y_i which shows that Z is m-connected. □

As Diestel, Jensen, Gorbunov and Thomassen showed, the existence of a large m-connected vertex set is an obstruction to the graph having low tree-width, thus the same holds for dexterous sets as well. Concretely, Diestel, Jensen, Gorbunov and Thomassen showed the following:

Proposition 3.4.21 ([28, Proposition 3(i)]). *Let G be a graph and $k > 0$ an integer. If G has tree-width $< k$ then G contains no $(k+1)$-connected set of size $\geqslant 3k$.*

Together with the grid theorem by Robertson and Seymour [67] (see also [28, Theorem 2]), which states that a graph of large enough tree-width needs to contain an $N \times N$-grid as a minor, this directly implies a quantitative relation between the existence of a dexterous set and a grid minor:

Theorem 3.4.22. *There is a function $f \colon \mathbb{N} \to \mathbb{N}$ such that every graph containing a dexterous set of size at least $f(N)$ also contains the $N \times N$-grid as a minor*

Proof. If a graph contains a dexterous set of size at least k, it has, by Proposition 3.4.21 and Proposition 3.4.20, tree-width at least $\frac{k}{3} - 1$. However, by the grid minor theorem [67] (see also [28, Theorem 2]) there is a function f such that every graph of tree-width $> f(N)$ contains an $N \times N$-grid as a minor. □

For large m-agile sets we will be able to show that their existence is, again quantitatively, characterized by $K_{m,k}$ and rectangular grid minors. Again we can build on existing literature about variations of the grid theorem. This time we will incorporate a generalized version of the grid theorem obtained from Geelen and Joeris ([50, Theorem 9.3]). To state this theorem, we need the following additional definition from that paper. Given parameters t, l, n a *(t, l, n)-wheel* is a graph obtained from a tree T with t vertices, a set Z of size l, a permutation $\pi : V(T) \to V(T)$ and a function $\psi : Z \to V(T)$ via the following construction: we start with n disjoint copies of T, called $T_1, \dots, T_n$. Let us denote the copy of $v \in V$ in T_i as v_i. We then add an edge between v_i and v_{i+1} for any vertex $v \in V$ and any index i between 1 and $n-1$. Then we add an edge between v_n and w_1 where $w = \pi(v)$. As a last step, for every $z \in Z$ and $z = \psi(z)$, we add an edge between z and every v_i.

A (θ, n)-*wheel* is any graph which is a (t, l, b)-wheel for some $t, l \in \mathbb{N}$ satisfying $2t + l = \theta$.

[50, Theorem 9.3] by Geelen and Joeris now implies the following:

Theorem 3.4.23 (see also [50, Theorem 9.3])**.** *There exists a function $f : \mathbb{N}^2 \to \mathbb{N}$ such that, given $\theta, n \in \mathbb{N}$ with $\theta \geqslant 2$ and $n \geqslant 3$, every graph G containing a θ-connected set U of size at least $f(\theta, n)$ contains a $K_{\theta,n}$ or a (θ, n)-wheel as a minor.*

Using this result we can now, again quantitatively, show that the existence of an m-agile set indeed is characterized by the existence of a large rectangular grid or a large complete bipartite graph as a minor. Concretely, we can show the following:

Theorem 8. *There is a function $f : \mathbb{N}^2 \to \mathbb{N}$ such that every graph containing an $((m-1)2m+1)$-agile set of size at least $f(m,k)$ contains $K_{m,k}$ or the $(2m-1) \times k$-grid as a minor.*

Proof. Every such $((m-1)2m+1)$-agile set is, by Proposition 3.4.20, also $((m-1)2m+1)$-connected, thus by Theorem 3.4.23 there is a function f such that every graph containing an $((m-1)2m+1)$-agile set of size at least $f(m,k)$ either contains $K_{(m-1)2m+1,k}$ or an $((m-1)2m+1,k)$-wheel as a minor. We are now going to show that such a wheel induces a $K_{m,k}$ or a $(2m-1) \times k$-grid as a minor. For this recall that such a wheel was constructed using a tree T of size t, say, and a set Z of central vertices of size z, say, so that $2t + z = (m-1)2m+1$.

Consider the tree T' obtained from T by adding every vertex in Z as a leaf to T in such a way that $z \in Z$ is adjacent to its neighbour $\psi(z)$ in T. Then T' has at least $(m-1)m+1$ many vertices. Thus, T' has at least m leafs or contains a path of length at least $2m+1$.

If T' contains a set L of m leafs, we can construct a $K_{m,k}$-minor in G as follows: for every leaf $v \in L$ of T', if v is a vertex of T we let X_v be the set of all $v_i \in T'$. If $v \in Z$ then $X_v = \{z\}$. Clearly every X_v is connected and the sets X_v will be the branch sets of the vertices of degree k in our $K_{m,k}$-minor. For the vertices of degree m we now take, for every $1 \leqslant i \leqslant k$, the rest of T_i, i.e. the set $X_i = T_i \setminus \bigcup_{v \in L} X_v$. Since every $v \in L$ is a leaf of T, the set X_i is connected in the wheel. Moreover, each X_i has a neighbour in X_v for every $v \in L$ and the X_i and X_v are all pairwise disjoint, which completes the construction of our $K_{m,k}$-minor.

If on the other hand T' contains a path of length at least $2m+1$, then T needs to contain a path P of length $2m-1$, as the vertices in Z were only added as leafs to T. This P directly corresponds to the $(2m-1)$-columns of a $(2m-1) \times k$-grid minor in G, i.e. the restriction of the wheel to the set of all those v_i for which $v \in P$, equals an $(2m-1) \times k$-grid, except for some additional edges. □

We remark that it can actually be shown that a $(t, l, k(\frac{t+l}{2} - 1))$-wheel itself induces a $\left(\frac{t}{2} + l\right)$-agile set of size k, for example by performing a pebble pushing argument (akin to the one used in [2]) on the tree T. Thus, we could as well have formulated Theorem 8 in terms of a corresponding wheel instead of a regular grid.

Let us end this section with one final observation regarding Question 3.4.1. While we have seen that the existence of a large 2-agile set alone is not enough to guarantee the existence of a $K_{2,k}$-minor, due to the regular strips, it turns out that, by requiring the existence of a large 3-agile set, we can actually guarantee the existence of a $K_{2,k}$-minor.

Theorem 3.4.24. *There exists a function $f\colon \mathbb{N} \to \mathbb{N}$ such that every graph containing a 3-agile set of size at least $f(k)$ also contains a $K_{2,k}$-minor.*

Proof. Let G be a graph containing a 3-agile set X of size $f(k) := N > 4$. Like in Lemma 3.4.3 we may assume that G is 2-connected: if (A, B) is a separation of G such that $A \cap B$ contains at most one vertex, then either A or B contains only one vertex of X, otherwise X would not be 3-agile. So suppose that A contains at most one vertex from X, we claim that $G' := G - (A \setminus B)$ is a subgraph of G which also contains a 3-agile set of size N. Indeed, it is easy to check $X \cap B \cup (A \cap B)$ is such a 3-agile set in G'.

Thus, by taking a subset-minimal subgraph of G containing a 3-agile set of size N, we may suppose that G is 2-connected.

Let T be a normal spanning-tree of G. Since $|G| \geqslant N$, by taking N large enough we can ensure that T either contains a vertex of degree at least $k+1$ or a path of length L, say. If T contains a vertex of degree at least $k+1$, it is easy to find the desired $K_{2,k}$-minor, as follows: if $t \in T$ has degree at least $k+1$, there are k distinct components of $G - t$ which do not contain the root. Now since T is normal and G is 2-connected, each of these components needs to send an additional edge to the path P in T between r and t. Thus, taking all the components as the branch sets of the degree-2 vertices, the vertex v itself as the branch set of one of the degree-k vertices and the vertices of $P - t$ as the branch set of the other degree-k vertex gives the desired $K_{2,k}$-minor.

So we may suppose that T does not contain a vertex of degree $k+1$, and thus contains a path P of length at least L. Let us suppose further that there are at least n vertices $v_1, \dots, v_n$ (enumerated starting from the root of T) on P such that there is a component C_i of $T - P$ with neighbour v_i in T (possibly the empty component) with the property that $C_i \cup \{v_i\}$ contains a vertex x_i in X. Note that we can indeed achieve the existence of such a path by taking a large enough N.

Consider some partition $A \dot\cup B \dot\cup C$ of X into three disjoint sets with the property that $x_i \in A$ whenever i is divisible by 3, $x_i \in B$ whenever i equals 1 modulo 3 and $x_i \in C$ whenever i equals 2 modulo 3. Since X is 3-agile, there are disjoint trees $T_0, T_1, T_2 \subseteq G$ such that $A \subseteq T_0$, $B \subseteq T_1$, and $C \subseteq T_2$. We may assume without loss of generality that every vertex on P belongs to either T_0 or T_1 or T_2. We say that a subpath P' of P is *T_i-free* if no vertex on P' is contained in T_i. Now if P would contain a long subpath P' which contains $v_j, \dots, v_{j+3k+2}$, say, and which is T_i-free for some i, then we find the desired $K_{2,k}$-minor: there are at least k of the vertices $v_j, \dots, v_{j+3k+2}$, for which the corresponding x_l lies in T_i. In particular, the corresponding components C_l are met by T_i. Since T is a normal spanning tree, there therefore exists, for every such component, a vertex on rTv_j which lies in T_i and is adjacent to that component. But now we obtain our desired $K_{2,k}$-minor by taking as one of the vertices of high degree the subpath P', as the other one the path rTv_j and as the vertices of degree 2 the components C_l mentioned above.

Hence, there cannot be a long T_i-free subpath, thus every long enough subpath of P meets all three T_i.

In particular, if we partition P into subpaths P_1, P_2,... P_{3k} each containing $3k+2$ of the v_l, then each of the P_j meets all the T_i. Thus, for each P_j there is a $0 \leqslant i \leqslant 2$ such that P_j contains a subpath on which all v_l are contained in T_i, put the preceding v_l and the successive v_l on P together meet both the other trees.

By the pigeonhole principle we find that for at least k of the P_j the chosen T_i is the same. But now we obtain a $K_{2,k}$-minor by taking the other two T_i's as the vertices of high degree, and use the paths P_j found above as the vertices of degree 2 in our $K_{2,k}$-minor. □

Chapter 4

Distinguishing

In this chapter we are concerned with tree-of-tangle theorems in various contexts. We start in Section 4.1 with a theorem unifying Theorem 1.1.3 as well as a version of Theorem 1.1.2: our splinter Lemma 10. We also demonstrate that this lemma allows us to develop tree-of-tangle theorems for a variety of other contexts, some of which could not be solved by Theorem 1.1.3 or Theorem 1.1.2. That section is joint work with Jakob Kneip and Maximilian Teegen and published in [39].

After that, we consider another application of Lemma 10 in Section 4.2: an application to separations of directed graphs. Such an application is not possible directly as the separations of a directed graph do not form a universe of separations. However, using an additional twist we can in fact use a variation of Lemma 10 to reprove a tree-of-tangles theorem for directed graphs by Giannopoulou, Kawarabayashi, Kreutzer and Kwon. We do so via defining an abstract notion, like abstract separation systems, which reflects the structure of separations of directed graphs and proving a corresponding tree-of-tangles theorem for a notion of tangles of these separations. The results from that section are yet unpublished and joint work with Maximilian Teegen, except for Lemma 12, which is joint work with Maximilian Teegen and Jakob Kneip.

In Section 4.3 we then turn back to abstract separation systems and find a strengthening of Theorem 1.1.3 in another direction, by proving a canonical version of Theorem 1.1.3, i.e. showing that the construction of the tree of tangles in Theorem 1.1.3 can actually be performed invariant under isomorphisms. This result is published in [36] and joint work with Jakob Kneip.

After having shown these various results about trees of tangles in finite structures, we turn our attention in Section 4.4 towards infinite separation systems. There, we develop a version of Lemma 10 for infinite separation systems and again show that this result can be used to obtain existing tree-of-tangles theorems for infinite graphs. While the results presented in this section are joint work with Maximilian Teegen, the paper [42] in which these results are published is joint work not just with Maximilian Teegen, but also with Jakob Kneip.

In Section 4.5 we then consider another application of this infinite splinter theorem: we show that it can be used to distinguish *edge-blocks* in a graph, where a k-edge-block is a $\subseteq$-maximal vertex set not separated by any cut of order less than k. This result can then be used to obtain an alternative proof of an important result by Dicks and Dunwoody, stating that there exists, in any graph, a nested set of cuts which 'generates' all cuts of that graph. This joint work with Jan Kurkofka and Maximilian Teegen is published in [43].

Turning our attention back to finite structures, the last part of this chapter is devoted not to a new tree-of-tangles theorem, but rather to a new proof strategy. While traditionally the tree-of-tangles theorem and the tangle-tree duality theorem form the two main independent pillars of tangle theory, we will be able to show that the second pillar – the tangle-tree duality theorem – can actually be used to obtain a proof of a tree-of-tangles theorem. Thus, these two pillars are not as independent, as one previously thought. Moreover, this proof allows us to give some bounds on the degrees of the nodes of the tree of tangles, which are not easily shown otherwise. These results are joint work with Jakob Kneip and Maximilian Teegen and published in [41].

4.1 Trees of tangles in abstract separation systems

4.1.1 Introduction

In this section of this thesis we bridge the gap between Theorem 1.1.2 and Theorem 1.1.3 by establishing the following tree-of-tangles theorem which combines some of the upsides of both Theorem 1.1.2 and Theorem 1.1.3, i.e. which is as widely applicable as Theorem 1.1.3 while still resulting in an efficient distinguisher as Theorem 1.1.2 when applied in a context with an order function:

Theorem 9. *If $\mathcal{S} = (S_1, \ldots, S_n)$ 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 $\mathcal{P}$.*

Theorem 9 includes Theorem 1.1.3 by taking a sequence of just one separation system. Also, if applied in the context of Theorem 1.1.2 we can obtain a nested set of separations efficiently distinguishing all the profiles, by taking as separation systems S_k the sets of all separations of order $< k$ of the given graph.

The nested set N found by Theorem 9 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 9, 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 10 (Splinter Lemma)**.** *Let U be a universe of separations and let $\mathfrak{A} = (\mathcal{A}_i)_{i \leqslant n}$ be 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 $\{a_1, \ldots, a_n\}$ is nested.*

Lemma 10, 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 10 not only implies Theorem 9, but can also be used to prove for example Theorem 1.1.1 and Theorem 1.1.3 directly. In fact Lemma 10 has a remarkably short proof (as we shall see in Section 4.1.2), making it the shortest available proof of Theorem 1.1.1 so far (see Section 4.1.3). Moreover, the premise in Lemma 10 is straightforward to check, and Lemma 10 itself does not make reference to tangles or any specific implementations of them. As a result Lemma 10 can be used in many different settings, implying variations of Theorem 1.1.1 in a multitude of contexts. For example, after deriving in Section 4.1.3 Theorem 1.1.1, Theorem 1.1.3, and Theorem 9 from Lemma 10, we use Lemma 10 to establish a new tree-of-tangles theorem in the setting of clique separations.

Since Lemma 10 does not yield a canonical set of separations, we cannot deduce the whole Theorem 1.1.2 from it. We fix this in Section 4.1.4 by establishing a version of Lemma 10 which does give a canonical nested set, albeit under slightly stronger assumptions:

Lemma 11 (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 we have that $N(\varphi(\mathfrak{A})) = \varphi(N(\mathfrak{A}))$.

We make use of Lemma 11 in Section 4.1.5 to obtain a new shortest proof of Theorem 1.1.2 and to extend Theorem 1.1.2 to two natural types of separations whose structural submodularity does not come from a submodular order function: clique separations, and circle separations.

4.1.2 The Splinter Lemma

In this section we establish Lemma 10, from which we shall derive two previously known theorems as well as two new flavours of tree-of-tangles theorems in Section 4.1.3. A cornerstone of the proofs of both Lemma 10 and of the two known results we shall derive from it is the fish Lemma 2.3.1.

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 2.3.1 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 [26].) 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 10 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 10 is met. Notably, Lemma 10 will make no reference at all to tangles or their specific properties. The proof of Lemma 10 will also utilize Lemma 2.3.1; 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 10. 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 $a_i \in \mathcal{A}_i \smallsetminus \mathcal{A}_j$ and

$a_j \in \mathcal{A}_j \smallsetminus \mathcal{A}_i$ 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 2.3.1 we are already able to state and prove our first main result:

Lemma 10 (Splinter Lemma)**.** *Let U be a universe of separations and let $\mathfrak{A} = (\mathcal{A}_i)_{i \leqslant n}$ be 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 $\{a_1, \ldots, a_n\}$ 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 2.3.1. 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 $\mathcal{A}_1, \ldots, \mathcal{A}_{n-1}$ to obtain a nested set consisting of some $a_1, \ldots, a_{n-1}$. 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 4.1.3 that this innocuous-looking lemma is actually strong enough to directly imply various existing tree-of-tangles theorems, including Theorem 1.1.1.

4.1.3 Applications of the Splinter Lemma

A short proof of Theorem 1.1.1

As a first application of Lemma 10 let us give a short proof of Theorem 1.1.1:

Theorem 1.1.1 ([68])**.** *Every graph has a tree-decomposition displaying its maximal tangles.*

If N is a nested set of separations of a finite graph G it is straightforward to find a tree-decomposition of G whose set of induced separations is precisely N(see [22, 68], this can also be shown by using the statements from Section 2.6). Therefore, in order to prove Theorem 1.1.1, it suffices to find a nested set N of separations of G which efficiently distinguishes all maximal tangles of G.

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

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

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

Let $\mathfrak{A}$ be the family of all these sets $\mathcal{A}_{P,P'}$. A nested set of separations of G distinguishes all maximal tangles of G efficiently if and only if it contains an element of each $\mathcal{A}_{P,P'}$. Therefore, the existence of such a set, and hence Theorem 1.1.1, now follows directly from Lemma 10 once we show that $\mathfrak{A}$ splinters:

Lemma 4.1.1. *The family $\mathfrak{A}$ of all $\mathcal{A}_{P,P'}$ splinters.*

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

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

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

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 P and P' if necessary we may assume that $(A, B) \in P$ and $(B, A) \in P'$. Then, by the above inequality, P 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 $(A, B) \in P$ and the tangle property (T). However, due to $(B, A) \in P'$ and the tangle property (T), P' 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 P and P' and hence lies in $\mathcal{A}_{P,P'}$. □

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 Theorem 1.1.3:

Theorem 1.1.3 ([26, Theorem 6])**.** *Let $\vec{S}$ be a structurally submodular separation system and $\mathcal{P}$ a set of profiles of S. Then S contains a nested set that distinguishes $\mathcal{P}$.*

The price to pay in Theorem 1.1.3 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 1.1.3 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 1.1.1, 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 1.1.3 from Lemma 10. 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 1.1.3 it suffices to show that the family

$$\mathfrak{A}_{\mathcal{P}} = (\mathcal{A}_{P,P'} \ : \ P \neq P' \in \mathcal{P})$$

splinters:

Lemma 4.1.2. *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 1.1.3 from Lemma 10.

Proof of Theorem 1.1.3. Let $\mathcal{P}$ be a set of profiles of S. By Lemma 4.1.2 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 10, 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 10 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 10 in the following sections.

Profiles in submodular universes

Theorem 1.1.3, which we deduced from Lemma 10 in the previous section implies that every graph has, for any fixed integer k, a tree-decomposition which displays its k-tangles. However, Robertson's and Seymour's Theorem 1.1.1 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 1.1.3 does not imply Theorem 1.1.1.

Moreover, since Theorem 1.1.3 does not assume that the universe U it is applied to comes with an order function, Theorem 1.1.3 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'.

Thus, let us recall Theorem 1.1.2, which satisfies both of the requirements above, and is the strongest tree-of-tangles theorem, in terms of its consequence, known so far:

Theorem 1.1.2 (Canonical tree-of-tangles theorem for separation universes [27, Theorem 3.6])**.** *Let $U = (\vec{U}, \leqslant, {}^*, \vee, \wedge, |\cdot|)$ 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.

As mentioned in Chapter 2, the definition of robustness of a set of profiles is rather involved and thus not repeated here. In the following proofs this definition will be used only in one place; therefore we shall use it there as a black box and refer the reader to [27] for the full definition. Alternatively, the reader may think of a 'robust set of profiles' as a slightly stronger 'distinguishable set of robust profiles'.

Since every k-tangle of a graph is robust ([27]), Theorem 1.1.2 indeed implies Theorem 1.1.1 of Robertson and Seymour that every graph has a tree-decomposition displaying its maximal tangles (see [27, Section 4.1] for more on building tree-decompositions from nested sets of separations, and how Theorem 1.1.2 implies Theorem 1.1.1). Moreover, Theorem 1.1.2 improves upon Theorem 1.1.1 by finding a *canonical* such tree-decomposition, i.e. one which is preserved by automorphisms of the graph. Since Lemma 10 does not guarantee any kind of canonicity, we are not able to deduce the full Theorem 1.1.2 from Lemma 10; however, using Lemma 10 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 1.1.2*. (In Section 4.1.4 we shall prove a version of Lemma 10 which implies Theorem 1.1.2 in full.)

Our strategy will largely be the same as in Section 4.1.3. 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 10 is the following:

Lemma 4.1.3. *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 1.1.2 from Lemma 10:

Proof of the non-canonical Theorem 1.1.2. By Lemma 4.1.3 the collection $\mathfrak{A}_{\mathcal{P}}$ of the sets $\mathcal{A}_{P,P'}$ splinters. Thus, by Lemma 10 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 $\mathcal{P}$. Since 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 profile in $\mathcal{P}$ is regular, which implies that T is a regular tree set in that case. □

Sequences of submodular separation systems

Let us, once more, compare Theorem 1.1.3 and Theorem 1.1.2. 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 9 which combines the advantages of both Theorem 1.1.3 and Theorem 1.1.2 (without canonicity), i.e. which is not dependent on the existence of some order function, but which is as powerful and efficient as Theorem 1.1.2 if such an order function does exist.

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

If $S_1 \subseteq S_2 \subseteq \ldots \subseteq S_n$ 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 $S_1 \subseteq S_2 \subseteq \ldots \subseteq S_n \subseteq U$ 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 4.1.3, i.e. if S_i is the set of all separations in U of order $< i$, then the sequence $S_1 \subseteq S_2 \subseteq \ldots \subseteq S_n \subseteq U$ is a compatible sequence of submodular separation systems.

A *profile in* $\mathcal{S} = (S_1, \ldots, S_n)$ 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 Chapter 2 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 [27]: 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 $(\vec{r} \vee \vec{s}) \in \vec{S}_j$.

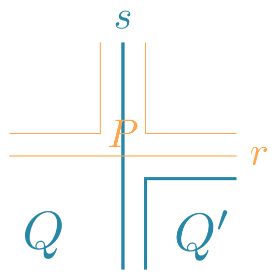

Figure 4.1: Robustness.

With the above definitions we are now able to formally state and prove Theorem 9, which includes both Theorem 1.1.3 and the non-canonical Theorem 1.1.2 (and hence Theorem 1.1.1) as special cases:

Theorem 9. *If $\mathcal{S} = (S_1, \ldots, S_n)$ 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 9 runs along very similar lines as the proof of Theorem 1.1.2 in the previous section, we only sketch 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 9 follows directly from Lemma 10 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'. □

Theorem 9 directly implies both Theorem 1.1.3 and the non-canonical Theorem 1.1.2: 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 $S_n = U$.

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 [55, 75]. More recently clique separations have received quite some attention in theoretical computer science (see for instance [1, 15, 62]) following Tarjan's work [71] on their algorithmic aspects.

In [26] 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 10 directly instead of Theorem 1.1.3, we are able to obtain a stronger result than the one given in [26], much in the same way that Theorem 9 improves upon Theorem 1.1.3.

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

It was shown in [26, Lemma 17] that S is a submodular separation system. Following their proof, we can show that in fact every $S_k \subseteq S$ is a submodular separation system, and that these extend each other in a way similar to the ordinary S_k of G:

Lemma 4.1.4. *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. □

We can now consider profiles in G with respect to these separation systems. A *profile P (of clique-separations) of order k* shall be a consistent orientation of S_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 such a profile P of order $|V|$ in G by letting P contain a separation $(A, B) \in \vec{S}$ 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 profile of order k in G. Let us denote by $\mathcal{P}_k$ the set of all profiles of order k.

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

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

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

Lemma 4.1.5. *For any set $\mathcal{P}$ of profiles the collection*

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

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 4.1.4 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 10 to obtain the following theorem:

Theorem 4.1.6. *There is a nested set of separations which efficiently distinguishes all the distinguishable profiles in* $\bigcup_{i=1}^{n} \mathcal{P}_i$.

Proof. By Lemma 4.1.5, we can apply Lemma 10 to

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

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 [22] for details). Thus, G admits a tree-decomposition whose adhesion sets (the sets $V_t \cap V_{t'}$ for distinct nodes t, t') 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 Tarjan in [71].

We will see in Section 4.1.5 that such a decomposition can in fact be chosen canonically, i.e. to be invariant under automorphisms of G.

4.1.4 Canonical Splinter Lemma

As we saw in the previous section, Lemma 10 is already strong enough to imply most of Theorem 1.1.2, but crucially does not guarantee the canonicity asserted in (iii). In this section we wish to prove a version of Lemma 10 using a stronger set of assumptions from which we can deduce Theorem 1.1.2 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 10 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 [27]. 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 10. 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:

(SH1) 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 $\mathcal{A}_i$.

(SH2) 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 $c_2 \in \mathcal{A}_i \cup \mathcal{A}_j$.

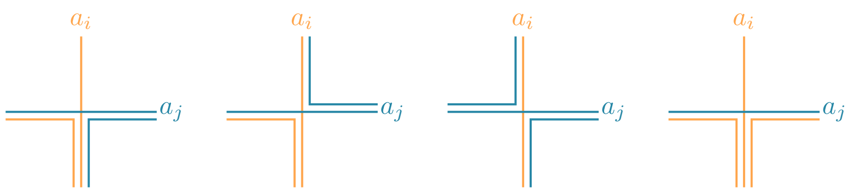

Figure 4.2: The possible configurations in (SH2) in the definition of *splinter hierarchically*, up to symmetry.

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 (SH2) 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 (SH1) for comparable elements of I is that otherwise we would not be able to deduce Theorem 1.1.2 in full from our main theorem of this section due to a quirk in the way that robustness is defined for profiles in [27] (see Section 4.1.5).

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 4.1.7. *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. Since a_i and a_j cross, there are three ways of orienting a_i and a_j such 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 (SH2) 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. □

We are now able to prove a canonical version of the splinter lemma by repeatedly applying Lemma 4.1.7 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 11 (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 we have that $N(\varphi(\mathfrak{A})) = \varphi(N(\mathfrak{A}))$.

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 4.1.7 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 4.1.7 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 2.3.1 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 $i \in I \smallsetminus J$. 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 (SH1), 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 2.3.1. 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 2.3.1. Therefore, (SH2) 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, applying the induction hypothesis yields $N(\varphi(\mathfrak{A}')) = \varphi(N(\mathfrak{A}'))$. 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. □

4.1.5 Applications of the Canonical Splinter Lemma

In this section we apply Lemma 11 to obtain a short proof of Theorem 1.1.2, to strengthen Theorem 4.1.6 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.

Robust profiles

Having established Lemma 11 in the previous section, we are now ready to derive the full version of Theorem 1.1.2. For this let $U = (\vec{U}, \leqslant, {}^*, \vee, \wedge, |\cdot|)$ 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 4.1.3, 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 1.1.2 from Lemma 11 as soon as we show that $\mathfrak{A}_{\mathcal{P}}$ splinters hierarchically.

Lemma 4.1.8. *$\mathfrak{A}_{\mathcal{P}}$ 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 (SH2) 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 and s have the same order, i.e. $|r| = |s|$; 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 (SH2) 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 (SH1) holds and concluding the proof. □

We are now ready to deduce the full Theorem 1.1.2 from Lemma 11:

Theorem 1.1.2 (Canonical tree-of-tangles theorem for separation universes [27, Theorem 3.6])**.** *Let $U = (\vec{U}, \leqslant, {}^*, \vee, \wedge, |\cdot|)$ 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.

Proof. By Lemma 4.1.8 the family $\mathfrak{A}_\mathcal{P}$ splinters hierarchically. Thus, we can apply Lemma 11 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 1.1.2.

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 11, 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. □

Clique separations

Regarding the profiles of clique separations discussed in Section 4.1.3, Lemma 4.1.4 not only suffices to show that the sets $\mathcal{A}_{P,P'}$ splinters, but can be used to show that the collection of these $\mathcal{A}_{P,P'}$ even splinters hierarchically, allowing us to apply Lemma 11: 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 4.1.4 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 $|r| < |s|$, (SH1) of splintering hierarchically is satisfied.

So suppose further that $|r| = |s|$, and let us check that (SH2) 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 (SH2) 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 4.1.4, 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 (SH2) is fulfilled.

Therefore, by Lemma 11 we get that we can choose the set in Theorem 4.1.6 canonically:

Theorem 4.1.9. *For every set $\mathcal{P}$ of profiles of clique separations of a graph G, there is a nested set $N = N(\mathcal{P})$ of separations which efficiently distinguishes all the distinguishable 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 11 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 $\mathcal{P}$. □

Circle separations

Another special case of separation systems are those of *circle separations* discussed in [26]: 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 subuniverse of the universe of all bipartitions of V:

Example 4.1.10. 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 $U = (\vec{U}, \leqslant, {}^*, \vee, \wedge, |\cdot|)$ the universe of all bipartitions of V with a submodular order function $|\cdot|$. Let $S \subseteq U$ be the set of all separations in U that are circle separations of V. Moreover, let us denote as S_k the set of all those circle separations in S whose order is $< k$.

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

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

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

Theorem 4.1.11. *The set S of all circle separations of V contains a tree set $T = T(S)$ that efficiently distinguishes all distinguishable tangles in 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 4.1.11 we need the following short lemma:

Lemma 4.1.12. *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 4.1.11.

Proof of Theorem 4.1.11. 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 11.

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 4.1.12, 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 (SH1) 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 (SH2) 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 (SH2) holds. □

4.2 Directed graphs

In this section we are going to see how we can reformulate our splinter Lemma 10 in order for us to make it applicable to distinguish tangles in directed graphs (digraphs). Such a theorem for digraphs has been developed by Giannopoulou, Kawarabayashi, Kreutzer, and Kwon in [51], where they showed that a directed tree-decomposition of a digraph which distinguish all its tangles of order k can be constructed in polynomial time. However, this tree-decomposition does not distinguish the tangles efficiently. The reason why constructing such a tree-decomposition in the setup of digraphs is more challenging lies in the definition of directed separations (see Section 4.2.1). These do not form a universe of separations, in particular one can not take corners of these directed separations in the same way as for separations of an undirected graph.

However, Giannopoulou, Kawarabayashi, Kreutzer, and Kwon also proved in [51] a result about efficient distinguishers: they were able to construct a structure called a 'tree-labelling' for the set of tangles (for definitions see Section 4.2.1). They then used this tree-labelling as an important step in the construction of the tree-decomposition mentioned above. More precisely, they showed the following:

Theorem 4.2.1 ([51, Theorem 6.2])**.** *Every set $\mathcal{T}$ of distinguishable tangles in a digraph G has a $\mathcal{T}$-tree-labelling.*

An important part of the proof of Theorem 4.2.1 is the case where $\mathcal{T}$ consists of distinct $(k+1)$-tangles which induce the same k-tangle. This case is then used to inductively construct the tree-labelling in Theorem 4.2.1. In this section of this thesis we show two results.

First, we are going to see that this important part of the proof of Theorem 4.2.1 can in fact be simplified by using a variation of our splinter Lemma 10; for this we are going to remark that Lemma 10 can in fact be formulated in an even more general context, where we do no longer require the existence of a separation system, but only a symmetric, reflexive relation corresponding to 'nestedness'. The same idea will also later be used in Section 4.4 to develop results like Lemma 10 for infinite structures.

This even more abstract version of Lemma 10 will then come in handy to achieve our second result of this section: using a different relation than before in this more abstract version of Lemma 10, we will be able to develop a theory of directed separation systems, and show that we can find a tangle distinguishing structure like the tree-labellings in Theorem 4.2.1 in this more abstract setting. Moreover, if applied to the directed separations of a digraph, this will give us a different proof of Theorem 4.2.1, solely based on this more abstract version of our splinter Lemma 10.

This Section 4.2 is structured as follows: first, in Section 4.2.1 we are going to give the definitions from [51] regarding digraphs and separations that we need. Then in Section 4.2.2, we are going to see how Lemma 10 can be generalized to work even without a separation system, and how this can be used to simplify the part of the proof of Theorem 4.2.1 mentioned above. After that, in Section 4.2.3, we are going to develop our theory of direct separation systems and use the more general version of our splinter lemma to prove a tree-of-tangles-like theorem for these directed separation systems. This yields an alternative proof of Theorem 4.2.1.

4.2.1 Directed graphs and separations

In this section we recall the definitions for separations and tangles in digraphs that we need. These definitions are largely based on [51], however, in some cases we are going to slightly adjust these definitions in order to make them more compatible with the definitions about separations and tangles used in other parts of this thesis, such as for graphs.

Given a digraph G and two vertex sets $A, B \subseteq V(G)$, a (directed) edge $e = (v, w) \in E(G)$ is said to *cross from A to B* if $v \in A \setminus B$ and $w \in B \setminus A$. An edge e is a *cross edge for A and B* if e crosses from A to B or from B to A.

Unlike in the undirected case, we allow for separations $\{A, B\}$ in digraphs the existence of edges between $A \setminus B$ and $B \setminus A$, as long as they all 'cross in the same direction'. With this definition a separation of a digraph reflects the notion of strong connectivity: recall that a vertex set in a digraph is *strongly connected* precisely if we find a directed path between any two vertices in that graph, i.e. there exist paths in 'both directions'. Thus, given a directed separation $\{A, B\}$, no strong component of $G - A \cap B$ will meet both A and B.

Formally, let us say that a pair (A, B) of vertex sets of G is an *(oriented) directed separation of G* if $A \cup B = V$ and there are either no cross edges from A to B or no cross edges from B to A. As for undirected graphs, the *separator* of (A, B) is $A \cap B$ and the *order* of (A, B) is the size of the separator, i.e. $|(A, B)| = |A \cap B|$. We call the (unoriented) pair $\{A, B\}$ the *underlying unoriented directed separation of* (A, B) (and (B, A)). In this section we will denote the set of all (unoriented) directed separations of a digraph G as $S(G)$.

Now given an oriented directed separation $X = (A, B)$, there are no edges that cross from A to B or no edges that cross from B to A. To distinguish these two cases, we also write the directed separation (A, B) as $X = (A \to B)$ if there are no cross edges from B to A, and as $X = (A \leftarrow B)$ if there are no cross edges from A to B. Moreover, if X is an unoriented directed separation, we denote as $X^+, X^- \in X$ the sides of X so that $(X^+ \to X^-)$ is an orientation of X, i.e. X^+ and X^- are chosen so that there are no edges from $X^- \setminus X^+$ to $X^+ \setminus X^-$.[1]

We can now define what a tangle in a digraph shall be. A *(directed) tangle of order k* in a digraph G is a set τ of oriented directed separations which all have order $< k$, such that for every unoriented directed separation $\{A, B\}$ of order $< k$ exactly one of (A, B) and (B, A) is contained in τ, and moreover any three separation $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$ satisfy the *tangle condition* in that $A_1 \cup A_2 \cup A_3 \neq V(G)$. We say that τ is a *tangle* if τ is a tangle of order k for some natural number k.

We say that an (unoriented) directed separation $\{A, B\} \in S(G)$ *distinguishes* two directed tangles τ, τ' if $(A, B) \in \tau$ and $(B, A) \in \tau'$ (for some orientation (A, B) of $\{A, B\}$). Two tangles are *distinguishable* if there exists a separation which distinguishes them. Otherwise, they are *indistinguishable*. Note that τ and τ' are indistinguishable precisely if $\tau \subseteq \tau'$ or $\tau' \subseteq \tau$. Given some integer l, two tangles τ and τ' are *l-distinguishable* if there exists a separation $\{A, B\}$ of order l which distinguishes τ and τ'. Similarly, we define two tangles to be *$< l$-distinguishable*, if there exists a separation $\{A, B\}$ of order less than l which distinguish them. If two tangles are not l-distinguishable, or not $< l$-distinguishable, they are *l-indistinguishable*, or *$< l$-indistinguishable*, respectively.

[1] If $X = \{A, B\}$ is a separation of the underlying undirected graph, i.e. there are no edges from $A \setminus B$ to $B \setminus A$ and no edges from $B \setminus A$ to $A \setminus B$, we choose X^+ and X^- arbitrary.

If $\{A, B\}$ is a separation of minimum possible order that distinguishes τ and τ', then $\{A, B\}$ distinguishes τ and τ' *efficiently*.

Finally, we define the tree-like structure distinguishing tangles obtained in [51]. Given a set $\mathcal{T}$ of tangles in a digraph G, let us say that a triple $\mathcal{L} := (L, \beta, \gamma)$ of an (undirected) tree L and two maps $\beta : V(L) \to \mathcal{T}$ and $\gamma : E(L) \to S(G)$ is a *tree-labelling of* $\mathcal{T}$ if $\mathcal{L}$ satisfies the following two properties:

(TL1) β is a bijection and

(TL2) for any two distinct nodes t, t' of L and the unique path P between them, the separation $\gamma(e)$ distinguishes $\beta(t)$ and $\beta(t')$ efficiently for every edge e of P, as long as e satisfies that $\gamma(e)$ has minimum possible order among all separations associated to P, i.e. $|\gamma(e)| = \min\{|\gamma(e')| : e' \in E(P)\}$.

Note that this definition implies that, given any edge e of T, the separation $\gamma(e)$ is a minimum order distinguisher for the images under β of the end vertices of e. In particular there are, for every $e \in E(L)$, tangles for which $\gamma(e)$ is an efficient distinguisher.

Now Theorem 4.2.1 states that, given any set $\mathcal{T}$ of pairwise distinguishable tangles, we can always find a tree-labelling for this set of tangles. A major step in the proof of Theorem 4.2.1 from [51] is the following intermediate result, which essentially allows one to prove Theorem 4.2.1 via an inductive argument:

Theorem 4.2.2 ([51, Theorem 6.3])**.** *Let $\mathcal{T}$ be a set of tangles of order $> l$ in a digraph G which are pairwise l-distinguishable but $< l$-indistinguishable. Then there is a $\mathcal{T}$-tree-labelling (L, β, γ) of G such that $|\gamma(e)| = l$ for all $e \in E(L)$.*

In the following, we are going to see how this result can actually be obtained from a corresponding version of our splinter Lemma 10. The challenge here lies in that, at least for the natural definition of declaring two unoriented directed separations $\{A, B\}$ and $\{C, D\}$ as *nested* if and only if they are nested as set separations, a tree-labelling does not give a nested set of separations. Moreover, with the natural partial order originating from set separations, they will also not form a universe of separations. Thus, we will not be able to use Lemma 10 as it is stated in Section 4.1 to prove Theorem 4.2.2. Instead, we are going to see in the next section that Lemma 10 can actually be formulated in a more general context with the exact same proof. This variation of Lemma 10 will then be applicable to obtain Theorem 4.2.2.

4.2.2 Abstracting Lemma 10 away from separation systems

As said above, the directed separations of a digraph do not form a universe of separations if equipped with the natural partial order for set separations as defined in Section 2.8. This is due to the fact that we can not guarantee that there are corners for any two oriented directed separations. In fact, we can only guarantee this for one of the two pairs of opposing corners: if given two directed separations X and Y then $(X^+ \cup Y^+, X^- \cap Y^-)$ and $(X^+ \cap Y^+, X^- \cup Y^-)$ will again give directed separations which will form one pair of opposite corners of X and Y, however neither $(X^+ \cup Y^-, X^- \cap Y^+)$ and $(X^+ \cap Y^-, X^- \cup Y^+)$ will in general be directed separations.

Hence, if we want to apply Lemma 10 in a directed setting, we would either need to come up with a different definition of a partial order on the directed separations, for which they do form a universe of separations, or we need to generalize the statement of Lemma 10 in order to be able to apply it to structures other than universes of separations. As it turns out, both the statement and proof of our splinter Lemma 10 can be formulated purely in terms of unoriented abstract separations and the nestedness relation between them. Consequently, we can define a more general structure encoding a set of unoriented separations together with a nestedness relation on them and can then obtain a version of Lemma 10 for these more general structures, i.e. a version of Lemma 10 which does not reference a universe of separations at all.

For this, let $\mathcal{A}$ be any finite set, and let $\sim$ be a reflexive and symmetric relation on $\mathcal{A}$. We will think of $\sim$ as the relation of 'being nested'; i.e. we will be able to obtain Lemma 10 from the more general version of the splinter lemma proved in this section as a simple corollary by applying it to our universe U of unoriented separations as our set $\mathcal{A}$ and taking as $\sim$ the nestedness relation on this set.

We now need to define what a 'corner' in this more general setup should be. Since a central ingredient of our proof of Lemma 10 is the fish Lemma 2.3.1, we will incorporate that lemma into our definition of corner. So let us say that given two elements $a, a' \in \mathcal{A}$, a *corner* of a and a' is any element $b \in \mathcal{A}$ with the property that $b \sim c$ for all $c \in \mathcal{A}$ satisfying both $a \sim c$ and $a' \sim c$. In other words, a corner of a and a' is an element of $\mathcal{A}$ that is nested with all those elements of $\mathcal{A}$ that are nested with both a and a'. In particular, unlike for the classic definition of a corner between two abstract separations in a universe, both a and a' themselves are valid corners of a and a'.

This definition is actually all we need in order to define what *splinters* should mean in our more abstract setup: given a finite set $\mathcal{A}$ together with a reflexive symmetric relation $\sim$ on $\mathcal{A}$, we say that a family $\mathfrak{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_n\}$ of subsets $\mathcal{A}_i$ of $\mathcal{A}$ *splinters* if, for any $a_i \in \mathcal{A}_i$ and any $a_j \in \mathcal{A}_j$, we have that there is a corner b of a_i and a_j which is either contained in $\mathcal{A}_i$ and satisfies $b \sim a_j$, or is contained in $\mathcal{A}_j$ and satisfies $b \sim a_i$. This means that for any two elements a_i and a_j from $\mathcal{A}_i$ and $\mathcal{A}_j$ there should be a corner of these two which is contained in one of the two sets $\mathcal{A}_i$ and $\mathcal{A}_j$ and is nested with the element a_j or a_i from the other of the two. Note that, if a_i and a_j are nested, i.e. $a_i \sim a_j$, then we can always take one of the two as the required corner, thus it is enough to verify the splinter condition for elements a_i and a_j which are not nested, i.e. which do not satisfy $a_i \sim a_j$.

Our more general version of Lemma 10 now reads as follows:

Lemma 12 (Abstract Splinter Lemma)**.** *Let $\mathcal{A}$ be a finite set and let $\sim$ be a reflexive and symmetric relation on $\mathcal{A}$. Let $\mathfrak{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_n\}$ be a family of subsets of $\mathcal{A}$ which splinters. Then there is a set $N = \{a_1, \dots, a_n\} \subseteq \mathcal{A}$ such that $a_i \in \mathcal{A}_i$ for every $1 \leqslant i \leqslant n$ and such that $a \sim a'$ for any $a, a' \in N$.*

We can prove Lemma 12 in exactly the same way as Lemma 10, for the reader's convenience we nevertheless include a proof here:

Proof of Lemma 12. 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 $1 \leqslant i \leqslant n$ and an $a_i \in \mathcal{A}_i$ so that a_i is nested with at least one element of $\mathcal{A}_j$ for each $j \neq i$, i.e. for every $j \neq i$ there is an $a_j \in \mathcal{A}_j$ such that $a_i \sim a_j$. Then the assertion holds: for $j \neq i$ let $\mathcal{A}'_j$ be the set of those elements a_j of $\mathcal{A}_j$ that are nested with a_i, i.e. satisfy $a_i \sim a_j$. Then $(\mathcal{A}'_j : j \neq i)$ is a family of non-empty sets which splinters by the definition of a corner: if, $a_j \in \mathcal{A}'_j$ and $a_k \in \mathcal{A}'_k$ for some $j, k \neq i$, then $a_i \sim a_j$ and $a_i \sim a_k$ and thus any corner b of a_j and a_k satisfies $b \sim a_i$. Hence, $(\mathcal{A}'_j : j \neq i)$ splinters because $(\mathcal{A}_i : 1 \leqslant i \leqslant n)$ did. Thus, by the induction hypothesis we can pick a set $N' = \{a_j \in \mathcal{A}'_j : j \neq i\}$ so that any two $a, b \in N$ satisfy $a \sim b$. Therefore, $N = N' \cup \{a_i\}$ is as required, as we have that $a \sim b$ for any two $a, b \in N$ and that N meets every $\mathcal{A}_i$.

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 set N' consisting of some $a_1, \dots, a_{n-1}$ which satisfy $a_j \sim a_k$ for all $1 \leqslant j, k \leqslant n-1$. Fix an arbitrary $a_n \in \mathcal{A}_n$. For all $i < n$, we have by our splinter condition that there either is a corner of a_n and a_i contained in $\mathcal{A}_n$ or that there is a corner contained in $\mathcal{A}_i$. If for some i we find a corner b between a_i and a_n that is contained in $\mathcal{A}_n$, this a_i is the desired separation for the above argument, as $a_i \sim a_j$ whenever $j \neq i$ and $j \neq n$ and moreover $a_i \sim b$ and $b \in \mathcal{A}_n$. Otherwise, for each $i < n$, there is a corner b_i between a_i and a_n contained in $\mathcal{A}_i$, thus we can apply the above argument to a_n, as $a_n \sim b_i$ for all $1 \leqslant i \leqslant n-1$. □

We now want to use Lemma 12 to obtain a proof of Theorem 4.2.2. For that let us define, given a set $\mathcal{T}$ of tangles of order $> l$ in a digraph G which are pairwise l-distinguishable but $< l$-indistinguishable, the set $\mathcal{A}$ as the set of all directed separations of order l and let us define, for $\tau, \tau' \in \mathcal{T}$ as $\mathcal{A}_{\tau,\tau'}$ the set of all those directed separations from $\mathcal{A}$ which efficiently distinguish τ and τ'.

We then want to apply Lemma 12 to the family $\mathfrak{A} = \{\mathcal{A}_{\tau,\tau'} : \tau, \tau' \in \mathcal{T}\}$. For this, we need to do define a relation $\sim$ on $\mathcal{A}$, in such a way that on the one hand, the family $\mathfrak{A}$ splinters with respect to $\sim$ and on the other hand, the set N of separations which we obtain from Lemma 12 allows us to obtain the desired tree-labelling for Theorem 4.2.2. As it turns out, the following relation will do exactly this.

Given a set $\mathcal{T}$ of tangles of a digraph, we say that two unoriented directed separations $\{A_1, B_1\}, \{A_2, B_2\}$ are *nested with respect to* $\mathcal{T}$ or $\mathcal{T}$*-nested*, denoted as $\{A_1, B_1\} \sim_{\mathcal{T}} \{A_2, B_2\}$, if the following holds:

(1) all tangles in $\mathcal{T}$ orient both $\{A_1, B_1\}, \{A_2, B_2\}$, and

(2) there are orientations (A_1, B_1) and (A_2, B_2) of $\{A_1, B_1\}, \{A_2, B_2\}$ such that either every tangle in $\mathcal{T}$ containing (A_1, B_1) also contains (A_2, B_2) or every tangle containing (A_2, B_2) also contains (A_1, B_1).

An equivalent definition can be obtained as follows: any unoriented directed separation $\{A, B\}$ which is oriented by all tangles in $\mathcal{T}$ gives a bipartition $\{\mathcal{T}_A, \mathcal{T}_B\}$ of $\mathcal{T}$ by defining $\mathcal{T}_A$ as the set of all tangles from $\mathcal{T}$ containing (B, A) and $\mathcal{T}_B$ as the set of all tangles from $\mathcal{T}$ containing (A, B). Now two unoriented directed separations $\{A_1, B_1\}$ and $\{A_2, B_2\}$ which are oriented by all tangles in $\mathcal{T}$ are nested with respect to $\mathcal{T}$ precisely if the two bipartitions of $\mathcal{T}$ induced by $\{A_1, B_1\}$ and $\{A_2, B_2\}$ are nested with respect to the usual definition of nestedness for bipartitions.

This relation of being $\mathcal{T}$-nested is clearly reflexive and symmetric. Let us now see that the family $\mathfrak{A} = \{\mathcal{A}_{\tau,\tau'} : \tau, \tau' \in \mathcal{T}\}$ defined above indeed splinters with respect to this relation. Note that the following lemma uses the definition of a corner in the setup of our nestedness relation.

Lemma 4.2.3. *Let $\mathcal{T}$ be a set of tangles of order $> l$ in a digraph G which are pairwise l-distinguishable but $< l$-indistinguishable. Let X_1 and X_2 be directed separations such that $|X_1| = |X_2| = l$. Suppose that X_1 distinguishes $\tau_1, \tau_1' \in \mathcal{T}$ whereas X_2 distinguishes $\tau_2, \tau_2' \in \mathcal{T}$. Then either X_1 and X_2 are nested with respect to $\mathcal{T}$ or there is a corner of X_1 and X_2 which has order l and distinguishes τ_1 and τ_1' or τ_2 and τ_2'.*

Proof. Let $X_1 = (A_1 \to B_1)$ and $X_2 = (A_2 \to B_2)$. Let us suppose that X and Y are not nested with respect to $\mathcal{T}$. Thus, there are tangles $\tau_{aa}, \tau_{ab}, \tau_{ba}, \tau_{bb}$ in $\mathcal{T}$ such that $(B_1, A_1), (B_2, A_2) \in \tau_{aa}$, $(A_1, B_1), (B_2, A_2) \in \tau_{ba}$, $(B_1, A_1), (A_2, B_2) \in \tau_{ab}$ and $(A_1, B_1), (A_2, B_2) \in \tau_{bb}$. Now both $(A_1 \cap A_2, B_1 \cup B_2)$ and $(A_1 \cup A_2, B_1 \cap B_2)$ are again directed separations. If one of the two, say $(A_1 \cap A_2, B_1 \cup B_2)$ would have order less than l, then, by the definition of a tangle, $(B_1 \cup B_2, A_1 \cap A_2) \in \tau_{aa}$, but $(A_1 \cap A_2, B_1 \cup B_2) \in \tau_{bb}$ contradicting the fact that τ_{aa} and τ_{bb} are $< l$-indistinguishable. Thus, both $(A_1 \cap A_2, B_1 \cup B_2)$ and $(A_1 \cup A_2, B_1 \cap B_2)$ have order exactly l.

Now the bipartition of $\mathcal{T}$ induced by $(A_1 \cap A_2, B_1 \cup B_2)$ is one of the four corners of the corresponding bipartitions of $\mathcal{T}$ induced by (A_1, B_1) and (A_2, B_2): a tangle $\tau \in \mathcal{T}$ contains $(A_1 \cap A_2, B_1 \cup B_2)$ if and only if the tangle contains one of (A_1, B_1) and (A_2, B_2). Thus, every separation Z which is $\mathcal{T}$-nested with X and Y is also $\mathcal{T}$-nested with both $(A_1 \cap A_2, B_1 \cup B_2)$ and $(A_1 \cup A_2, B_1 \cap B_2)$.

Now let us assume that one of $(A_1 \cap A_2, B_1 \cup B_2)$ and $(A_1 \cup A_2, B_1 \cap B_2)$ also distinguishes τ_1 and τ_1', say $(A_1 \cap A_2, B_1 \cup B_2)$ does so. Then $(A_1 \cap A_2, B_1 \cup B_2)$ and (A_2, B_2) are nested with respect to $\mathcal{T}$ and thus $(A_1 \cap A_2, B_1 \cup B_2)$ is the desired corner.

If on the other hand neither $(A_1 \cap A_2, B_1 \cup B_2)$ nor $(A_1 \cup A_2, B_1 \cap B_2)$ distinguish τ_1 and τ_1' then both τ_1 and τ_1' need to contain $(A_1 \cap A_2, B_1 \cup B_2)$ and $(B_1 \cap B_2, A_1 \cup A_2)$. Thus, by the definition of a tangle, neither τ_1 nor τ_1' can contain both (A_1, B_1) and (A_2, B_2) or both, (B_1, A_1) and (B_2, A_2). Thus, as (A_1, B_1) distinguishes τ_1 and τ_1', one of the two, say τ_1, needs to contain (A_1, B_1) and (B_2, A_2) whereas the other one needs to contain (B_1, A_1) and (A_2, B_2). Thus, (A_2, B_2) distinguishes τ_1 and τ_1'. But this implies that (A_2, B_2) is the desired corner. □

We now have all the ingredients needed to use Lemma 12 to prove Theorem 4.2.2.

Proof of Theorem 4.2.2. For any two tangles $\tau_1 \neq \tau_2 \in \mathcal{T}$, let $\mathcal{A}_{\tau_1,\tau_2}$ be the set of all unoriented directed separations which efficiently distinguish τ_1 and τ_2, and let $\mathcal{A} = \bigcup_{\tau_1,\tau_2 \in \mathcal{T}} \mathcal{A}_{\tau_1,\tau_2}$

By Lemma 4.2.3 the family $\mathfrak{A}$ of all these sets $\mathcal{A}_{\tau_1,\tau_2}$ splinters with respect to the relation $\sim_{\mathcal{T}}$ on $\mathcal{A}$ of being nested with respect to $\mathcal{T}$. Thus, by Lemma 12, there is a set N of directed separations which meets all the $\mathcal{A}_{\tau_1,\tau_2}$ such that any two separations from N are $\mathcal{T}$-nested. Let us now see how we can use this set N to obtain a tree-labelling for $\mathcal{T}$.

For this let us assume that N is $\subseteq$-minimal with the property that N meets every $\mathcal{A}_i$. Let us take $\mathcal{T}$ as the vertex set of our tree L, and consequently let the map β be the identity. We add an edge e between τ and τ' if and only if $|N \cap \mathcal{A}_{\tau,\tau'}| = 1$. In that case, we let $\gamma(e)$ be the unique separation in $N \cap \mathcal{A}_{\tau,\tau'}$. Additionally, we shall define a map ρ from the set of orientations of the edges of L, i.e. from the set $\vec{E}(L) := \{(v,w) \,:\, \{v,w\} \in E(L)\}$ to the set of oriented separations in N, which maps the edge (τ,τ') to the orientation of the separation in $N \cap \mathcal{A}_{\tau,\tau'}$ contained in τ'.

Let $\vec{\mathcal{B}}$ be the set of bipartitions of $\mathcal{T}$ induced by the directed separations in N, i.e. for every directed separation (A,B) with $\{A,B\} \in N$, we add to $\vec{\mathcal{B}}$ the bipartition $(A_\tau, B_\tau) = (\{\tau \in \mathcal{T} \,:\, (B,A) \in \tau\}, \{\tau \in \mathcal{T} \,:\, (A,B) \in \tau\})$. Since N is a set consisting of separations which are pairwise $\mathcal{T}$-nested, the set $\vec{\mathcal{B}}$ consists of bipartitions which are pairwise nested for the usual definition of nestedness of bipartitions. Moreover, $\mathcal{T}$ naturally induces a set $\mathcal{O}$ of consistent orientations of $\vec{\mathcal{B}}$: given $\tau \in \mathcal{T}$ we orient every bipartition in $\vec{\mathcal{B}}$ towards the side containing τ. Since N is a $\subseteq$-minimal set distinguishing all the tangles in $\mathcal{T}$, the set $\vec{\mathcal{B}}$ is $\subseteq$-minimal with the property that it distinguishes all the orientations in $\mathcal{O}$ for the usual definitions for bipartitions.

Let (T,α) be the order-respecting $\mathcal{B}$-tree (T,α) from Corollary 2.6.5. Then $V(T)$ equals $\mathcal{O}$ and, given a directed edge $(O,O') \in \vec{E}(T)$, we have that $\alpha(O,O')$ is the unique bipartition in $\vec{\mathcal{B}}$ which distinguishes O and O'. Thus, L is isomorphic to (T,α) via the isomorphism μ mapping a tangle in $\mathcal{T}$ to the orientation in $\mathcal{O}$ it induces. Moreover, given any directed edge (τ,τ') from L, the separation $\rho(\tau,\tau')$ induces the bipartition $\alpha(\mu(\tau),\mu(\tau'))$.

However, this isomorphism clearly implies that L is a tree and that our construction indeed gives a tree-labelling: given any two tangles τ,τ' and an edge e on the path between them, we can consider the corresponding edge e' on the path between $\mu(\tau)$ and $\mu(\tau')$ in T and know that the bipartition $\alpha(\vec{e'})$ distinguishes $\mu(\tau)$ and $\mu(\tau')$. Consequently, so does the directed separation $\rho(\vec{e})$ and consequently, the undirected separation $\gamma(e)$ corresponding to $\rho(\vec{e})$ does so as well. $\square$

In the next section we are going to develop an abstract version of these directed separations and are going to see that the ideas of this proof will actually allow us to obtain a version of Theorem 4.2.1 in this even more abstract setting.

4.2.3 Abstract directed separation systems

As said in the previous sections, the separations of a digraph are not covered from our usual notions of abstract separation systems inside a universe of separations. This is why, in this section, we are going to develop a directed generalization of abstract separation systems which does cover the separations of a digraph as an instance. We prove a corresponding version of a tree-of-tangles-type theorem for these directed analogue which in particular allows us to obtain a proof of Theorem 4.2.1 as a corollary.

Unlike the separations of an undirected graph, the separations $\{A,B\}$ of a digraph come with a natural 'direction' $(A \to B)$. To encode this in our abstract theory, our directed separation system will be the union of two sets, which in the instance of directed separations $\{A,B\}$ of a digraph correspond to the set of all the oriented separations of the form $(A \to B)$ and the set of all the oriented separations of the form $(B \leftarrow A)$. So let us say that a *directed separation system* is a poset $\overleftrightarrow{S} = \vec{S} \cup \overleftarrow{S}$ together with an order reversing involution $^* : \overleftrightarrow{S} \to \overleftrightarrow{S}$ such that, for $\vec{s} \in \vec{S}$ we have $\overleftarrow{s} := \vec{s}^* \in \overleftarrow{S}$, and vice versa. Note that we do not require the two sets $\vec{S}$ and $\overleftarrow{S}$ do be disjoint, nevertheless it is easy to see that the involution * gives a bijection between $\vec{S}$ and $\overleftarrow{S}$. We will usually denote elements of $\overleftrightarrow{S}$ as $\overleftrightarrow{s}$ if we do not want to specify whether they lie in $\vec{S}$ or $\overleftarrow{S}$.

Given such a directed separation system, we denote, given a separation $\vec{s} \in \vec{S}$ the pair $\{\vec{s}, \overleftarrow{s}\}$ just as s and the set of all these pairs as S. The elements of S are then also called *unoriented directed separations*. Conversely, given a subset N of S we denote as $\vec{N}$ the set of all $\vec{s} \in \vec{S}$ such that $s \in N$. Analogously we define $\overleftarrow{N}$ and $\overleftrightarrow{N}$.

If a directed separation system $\overleftrightarrow{U}$ is such that $\vec{U}$ and $\overleftarrow{U}$ are both lattices, then we say that $\overleftrightarrow{U}$ (or U) is a *directed universe of (directed) separations*.

Given a directed separation system $\overleftrightarrow{S}$, an *orientation of* $\overleftrightarrow{S}$ is a set $O \subseteq \overleftrightarrow{S}$ such that for every $\vec{s} \in \vec{S}$ we have that $|O \cap \{\vec{s}, \overleftarrow{s}\}| = 1$.

Such an orientation O is *consistent* if there do not exist any two distinct separations $\overleftrightarrow{s}, \overleftrightarrow{t} \in O$ such that $\overleftrightarrow{s}^* \leqslant \overleftrightarrow{t}$.

Given a directed universe $\overleftrightarrow{U}$, some separation system $\overleftrightarrow{S} \subseteq \overleftrightarrow{U}$, and some consistent orientation P of $\overleftrightarrow{S}$, we say that P is a *profile of* $\overleftrightarrow{S}$ 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 separation $\overleftrightarrow{s} \in S$ is *small* if $\overleftrightarrow{s} \leqslant \overleftrightarrow{s}^*$. A profile P is *regular* if there is no separation $\overleftrightarrow{s} \in P$ such that $\overleftrightarrow{s}^*$ is small.

A natural notion of structural submodularity for these directed separation systems now is the following: a directed separation system $\overleftrightarrow{S} \subseteq \overleftrightarrow{U}$ is *structurally submodular* if for any two $\vec{r}, \vec{s} \in \vec{S}$, we have that $\vec{r} \vee \vec{s} \in \vec{S}$ or $\vec{r} \wedge \vec{s} \in \vec{S}$. Note that this implies that also, for any $\overleftarrow{r}, \overleftarrow{s} \in \overleftarrow{S}$, we have that $\overleftarrow{r} \vee \overleftarrow{s} \in \overleftarrow{S}$ or $\overleftarrow{r} \wedge \overleftarrow{s} \in \overleftarrow{S}$, since * is an order reversing involution.

Similarly, there is a natural definition of a submodular order function: a function $f : \overleftrightarrow{U} \to \mathbb{N}$ is called an *order function* if f symmetric in that $f(\vec{s}) = f(\overleftarrow{s})$ for all $\vec{s} \in \vec{U}$. It is called *submodular* if $f(\vec{s}) + f(\vec{t}) \geqslant f(\vec{s} \vee \vec{t}) + f(\vec{s} \wedge \vec{t})$ for any two $\vec{s}, \vec{t} \in \vec{U}$. Note that this together with the fact that f is symmetric implies that $f(\overleftarrow{s}) + f(\overleftarrow{t}) \geqslant f(\overleftarrow{s} \vee \overleftarrow{t}) + f(\overleftarrow{s} \wedge \overleftarrow{t})$ for all $\overleftarrow{s}, \overleftarrow{t} \in \overleftarrow{U}$.

Given a submodular order function on a directed universe $\overleftrightarrow{U}$, the set $\overleftrightarrow{S}_k = \{\overleftrightarrow{s} \in \overleftrightarrow{U} \ : \ f(\overleftrightarrow{s}) < k\}$ is a structurally submodular directed separation system. In that case we say that a profile P of $\overleftrightarrow{S}_k$ is a *k-profile in* U, and that P is a *profile in* U if P is a k-profile for some integer k.

We can now, given a set of profiles $\mathcal{P}$ of a directed separation system $\overleftrightarrow{S}$, generalize the above defined notion of $\mathcal{P}$-nestedness for separations of a digraph to arbitrary directed separation systems: we say that two unoriented directed separations s and t from S are *$\mathcal{P}$-nested* if there are orientation $\overleftrightarrow{s}$ and $\overleftrightarrow{t}$ of s and t such that either every profile in $\mathcal{P}$ containing $\overleftrightarrow{s}$ also contains $\overleftrightarrow{t}$, or every profile in $\mathcal{P}$ containing $\overleftrightarrow{t}$ also contains $\overleftrightarrow{s}$. A set of separations is called *$\mathcal{P}$-nested* if any pair of separations from that set is $\mathcal{P}$-nested. We remark that this definition, just like the corresponding definition of $\mathcal{P}$-nestedness for separations of a digraph, can also be formulated in terms of the induced bipartitions of $\mathcal{P}$: given an oriented directed separation $\overleftrightarrow{s} \in \overleftrightarrow{S}$, this separation naturally induces a bipartition of $\mathcal{P}$ into the two sets $\mathcal{P}_{\vec{s}} := \{P \in \mathcal{P} \ : \ \vec{s} \in P\}$ of those profiles containing $\vec{s}$ and $\mathcal{P}_{\overleftarrow{s}} := \{P \in \mathcal{P} \ : \ \overleftarrow{s} \in P\}$ of those profiles containing $\overleftarrow{s}$. Now two unoriented directed separations are nested precisely if the bipartitions of $\mathcal{P}$ induced by their orientations are nested for the usual notion of nestedness of bipartitions.

Given two profiles P and P' of potentially distinct directed separation systems inside the same universe $\overleftrightarrow{U}$, a separation $s \in U$ *distinguishes* P, P' if there exists an orientation $\overleftrightarrow{s}$ of s such that $\overleftrightarrow{s} \neq \overleftrightarrow{s}^*$, and $\overleftrightarrow{s} \in P$ but $\overleftrightarrow{s}^* \in P'$. If $\overleftrightarrow{U}$ happens to be a universe with a submodular order function, then *s distinguishes P and P' efficiently* if s has minimum possible order among all separations in U distinguishing P and P'.

An example of a universe U of directed separations is the set $S(G)$ of directed separations of a digraph as defined in Section 4.2.1. This can be seen by defining $\vec{U}$ as the set of all the directed separations $(A \to B)$ of G and $\overleftarrow{U}$ as the set of all the directed separations $(A \leftarrow B)$ of G[2], and declaring $(A, B) \leqslant (C, D)$ if and only if $A \subseteq C$ and $D \supseteq B$. The involution on U is given by $(A \to B)^* = (B \leftarrow A)$ and $(B \leftarrow A)^* = (A \to B)$, for any separation $(A \to B) \in \vec{U}$ and any separation $(B \leftarrow A) \in \overleftarrow{U}$. Moreover, both $\vec{U}$ and $\overleftarrow{U}$ are indeed lattices which can be seen from the fact that, given two separations $(A \to B), (C \to D) \in \vec{U}$, both $((A \cup C) \to (B \cap D))$ and $((A \cap C) \to (B \cup D))$ are again directed separations. In this setting a separation (A, B) is small precisely if $B = V(G)$.

A natural submodular order function on this universe $U = S(G)$ is given by defining $|(A, B)| := |A \cap B|$. Consequently, if we let $\overleftrightarrow{S}_k(G)$ be the set of all directed separations of order at most k in $\overleftrightarrow{U}$, this gives us a structurally submodular directed separation system inside $\overleftrightarrow{U}$. Moreover, it is easy to see that if τ is a k-tangle of G as in the definition above, i.e. τ does not contain any three separations $(A_1, B_1), (A_2, B_2), (A_3, B_3)$ such that $A_1 \cup A_2 \cup A_3 = V(G)$, then τ is a regular k-profile in U.

Another example of a directed separation system is given by an ordinary abstract separation system $\vec{S}$. If we set $\overleftarrow{S} = \vec{S}$ and define the involution on $\overleftrightarrow{S} := \vec{S} \cup \overleftarrow{S} = \vec{S}$ the same as the involution on the abstract separation system $\vec{S}$, then this $\overleftrightarrow{S}$ forms a directed separation system. Moreover, with this translation also the concept of universes, profiles and (structurally) submodularity translates from the classical notions for abstract separation systems to the corresponding notions for our newly defined directed separation systems. In that sense, these new definitions of directed separation systems and universes generalize the existing theory of abstract separation systems.

We are now ready to use the exact same proof strategy as in our proof of Theorem 4.2.2 using Lemma 12 to obtain a similar statement in the setting of a directed separation system:

Theorem 4.2.4. *Let U be a directed universe with a submodular order function, let $k, l \in \mathbb{N}$ such that $l < k$ and let $\mathcal{P}$ be a set of profiles of S_k such that any two profiles in $\mathcal{P}$ are efficiently distinguished by a separation of order l. Then there exists a set $N \subseteq S$ of separations which is $\mathcal{P}$-nested and distinguishes any two profiles in $\mathcal{P}$ efficiently.*

Proof Sketch. For distinct profiles $P, P' \in \mathcal{P}$, let $\mathcal{A}_{P,P'}$ be the set of all separations from S_k efficiently distinguishing P and P' and note that these separations all have order exactly l.

As in the proof of Lemma 4.2.3, we can show that the family $\mathfrak{A}$ of all theses $\mathcal{A}_{P,P'}$ splinters. For this suppose we are given $s \in \mathcal{A}_{P,P'}$ and $t \in \mathcal{A}_{Q,Q'}$ which

[2]if (A, B) is a separation of the underlying undirected graph, i.e. there are no cross edges from $A \setminus B$ to $B \setminus A$ and no cross edges from $B \setminus A$ to $A \setminus B$, then (A, B) shall be contained in both, $\vec{U}$ and $\overleftarrow{U}$, which is covered by our definition, as we do not require the two sets $\vec{U}$ and $\overleftarrow{U}$ to be disjoint

are not $\mathcal{P}$-nested. If either $s \in \mathcal{A}_{Q,Q'}$ or $t \in \mathcal{A}_{P,P'}$ then s or t is the desired corner separation of s and t. Otherwise, either one of P and P' contains both $\vec{s}$ and $\vec{t}$, or one of them contains both $\overleftarrow{s}$ and $\overleftarrow{t}$. Moreover, as s and t are not $\mathcal{P}$-nested, there are profiles P_1 and P_2 in $\mathcal{P}$ such that $\vec{s}, \overleftarrow{t} \in P_1$ and $\overleftarrow{s}, \vec{t} \in P_2$. Since P_1 and P_2 cannot be distinguished by a separation of order less than l, this implies, by submodularity, that both $\vec{s} \vee \vec{t}$ and $\vec{s} \wedge \vec{t}$ need to have order l. But then one of these two separations is the corner required for the splinter condition.

Thus, by Lemma 12, there is a set $N \subseteq S$ with the desired properties. □

We remark that the exact same proof would actually allow us to obtain a slightly stronger result. Namely, we could in the setup of Theorem 4.2.4 define a set $N \subseteq S$ of separations to be *strongly $\mathcal{P}$-nested* if for every two separations $\vec{s}, \vec{t} \in \vec{N}$ either $\vec{s} \leqslant \vec{t}$, or $\vec{t} \leqslant \vec{s}$, or there are no profiles $P, P' \in \mathcal{P}$ such that $\vec{s}, \vec{t} \in P$ and $\overleftarrow{s}, \overleftarrow{t} \in P'$. We note that this property commutes with * and hence that we also have for any $\overleftarrow{s}, \overleftarrow{t} \in \overleftarrow{N}$ that either $\overleftarrow{s} \leqslant \overleftarrow{t}$, or $\overleftarrow{t} \leqslant \overleftarrow{s}$ or there does not exist profiles $P, P' \in \mathcal{P}$ such that $\overleftarrow{s}, \overleftarrow{t} \in P$ and $\vec{s}, \vec{t} \in P'$.

Moreover, if s and t are strongly $\mathcal{P}$-nested, then s and t are also $\mathcal{P}$-nested. But conversely, if s and t are $\mathcal{P}$-nested, this may be the case just because there is no profile in $\mathcal{P}$ containing $\vec{s}$ and $\overleftarrow{t}$. Thus, it is possible that s and t are not strongly $\mathcal{P}$-nested as we do not find orientations $\vec{s}, \vec{t}$ of s and t such that $\vec{s} \leqslant \vec{t}$ or $\vec{t} \leqslant \vec{s}$ and there are profiles $P, P' \in \mathcal{P}$ such that $\vec{s}, \vec{t} \in P$ and $\overleftarrow{s}, \overleftarrow{t} \in P'$.

The proof of Theorem 4.2.4 actually allows us to deduce that there always exists a strongly $\mathcal{P}$-nested set of separations efficiently distinguishing any two profiles in $\mathcal{P}$ in the above setup.

However, the property of being strongly $\mathcal{P}$-nested is only slightly stronger than the property of being $\mathcal{P}$-nested. Since we also do not have a use case for this stronger theorem, we preferred to work with the more natural notion of $\mathcal{P}$-nestedness, since that relation corresponds nicely to the nestedness relation of the induced bipartitions of $\mathcal{P}$.

Now it would be nice if we could show that, for a given set of profiles $\mathcal{P}$ of a structurally submodular directed separation system, we can always find a $\mathcal{P}$-nested set of separations which distinguish all the profiles in $\mathcal{P}$. However, as we will see in a moment, it is in generally not possible to find such a set. A necessity to find such a set would be a positive answer to the following question:

Question 4.2.5. Let U be a directed universe, let $S \subseteq U$ be structurally submodular and let $\mathcal{P}$ be a set of profiles of S. Does there exist a separation $s \in S$ which distinguishes some pair of profiles in $\mathcal{P}$ such that, for any two profiles $P, Q \in \mathcal{P}$ which are not distinguished by s, there exists a separation $t \in S$ which distinguishes P and Q and is $\mathcal{P}$-nested with s?

Moreover, actually any statement about distinguishing the profiles of a structurally submodular directed separation system in some sort of 'tree-like' way the author could think of would imply a positive answer to Question 4.2.5.

But unfortunately, the general answer to Question 4.2.5 is 'no', as we shall see in Example 4.2.6. Thus, it seems like the notion of directed structurally submodularity is too weak to be able to distinguish profiles in some sort of 'tree-like' way.

Example 4.2.6. Consider the system of directed bipartitions of a set of 8 points depicted in Fig. 4.3:

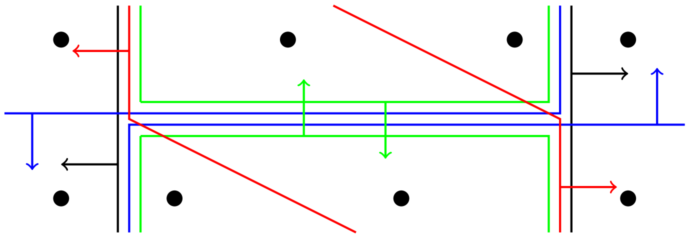

Figure 4.3: Each directed separation is represented by a coloured line, and the arrow on that line represents the direction for which that separation is contained in $\vec{S}$.

It is easy to calculate that this system is indeed a structurally submodular directed separation system.

Moreover, each of the 8 points of the ground set induces a directed profile of the separation system, let us denote as $\mathcal{P}$ the set of these 8 profiles. Now, for each separation in S there is a pair of profiles in $\mathcal{P}$ for which that separation is the only one distinguishing that pair of profiles. As moreover no separation in S is $\mathcal{P}$-nested with all the other separations in S, this example shows that we cannot find a separation s as in Question 4.2.5.

In light of this example, it seems difficult to obtain any useful theorems about structurally submodular directed separation systems. Therefore, we are instead going to work for the rest of this section with directed separation systems S_k induced by a submodular order function.

However, it will not be necessary for us to restrict ourselves to a set of profiles $\mathcal{P}$ of one fixed such S_k. Instead, we will be able to work with profiles of some universe U with a submodular order function, i.e. we do not require our profiles to all be k-profiles for the same k. Thus, while we are, in a way, not able to obtain a directed analogue of Theorem 1.1.3, we can prove, in some sense, a non-canonical directed analogue of Theorem 1.1.2.

For the rest of this section suppose that we are given some directed universe $\overleftrightarrow{U}$ with a submodular order function. Since we want to work with profiles of different orders, we need to come up with a variation of our notion of $\mathcal{P}$-nestedness which can deal with a set $\mathcal{P}$ containing profiles of different order. Recall that $\mathcal{P}$-nestedness is only defined for separations which are oriented by all profiles in $\mathcal{P}$.

Luckily, it is possible to define a notion of nestedness with respect to a set $\mathcal{P}$ of profiles of different orders, such that the sets $\mathcal{A}_{P,P'}$ of separations efficiently distinguishing the pair P, P' of profiles from $\mathcal{P}$ will indeed splinter with respect to that relation. Moreover, this relation will still be strong enough to allow us to use, when applied to the directed separations of a digraph, the set N from our splinter lemma to obtain the tree-labelling from Theorem 4.2.1.

So, for the rest of this section, let $\mathcal{P}$ be a set of pairwise distinguishable profiles in our directed universe $\overleftrightarrow{U}$ and note that the profiles in $\mathcal{P}$ do not need to all be profiles of the same separation system $\overleftrightarrow{S}_k$.

Let us define when two separations s and t from U are *weakly $\mathcal{P}$-nested*. This definition will make a case distinction depending on the relation between the order of s and t. If $|s| < |t|$, then we say that s and t are *weakly $\mathcal{P}$-nested* if there exists an orientation $\overleftrightarrow{s}$ of s such that any two profiles P and Q from $\mathcal{P}$ which are $< |s|$-indistinguishable and are distinguished by t do not both contain $\overleftrightarrow{s}$.

If $|s| = |t|$, then we say that s and t are *weakly $\mathcal{P}$-nested* if there are orientations $\overleftrightarrow{s}$ and $\overleftrightarrow{t}$ of s and t such that there is no pair of profiles P and Q from $\mathcal{P}$ which are $< |s|$-indistinguishable and distinguished by s where P and Q both contain $\overleftrightarrow{t}$ and, analogously, there is no pair of profiles P and Q from $\mathcal{P}$ which are $< |s|$-indistinguishable and distinguished by t where P and Q both contain $\overleftrightarrow{s}$.

A set of separations is *weakly $\mathcal{P}$-nested* if every pair of separations from that set is weakly $\mathcal{P}$-nested.

We observe that if all profiles in $\mathcal{P}$ are profiles of the same separation system S_k, then any set of separations of order less than k which is nested with respect to $\mathcal{P}$ is also weakly $\mathcal{P}$-nested. However, not every pair s, t of separations which is weakly $\mathcal{P}$-nested with respect to such a set $\mathcal{P}$ of profiles need to be $\mathcal{P}$-nested, as the definition of weak $\mathcal{P}$-nestdness of s and t only takes those pairs of profiles from $\mathcal{P}$ into account which are $< \min\{|s|, |t|\}$-indistinguishable.

We now want to prove the following:

Theorem 13. *If $\mathcal{P}$ is a set of distinguishable profiles in a directed universe $\overleftrightarrow{U}$, then there exists a weakly $\mathcal{P}$-nested set of separations which efficiently distinguishes every two profiles from $\mathcal{P}$.*

For this, we first need a lemma which gives us a characterization of when two separations of the same order are not weakly $\mathcal{P}$-nested:

Lemma 4.2.7. *Given two separations $s, t \in U$, if $|s| = |t|$, then s and t are not weakly $\mathcal{P}$-nested precisely if there are four profiles $P_{\vec{s},\vec{t}}, P_{\vec{s},\overleftarrow{t}}, P_{\overleftarrow{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}}$ in $\mathcal{P}$ which are pairwise $< |s|$ indistinguishable and which satisfy $\overleftrightarrow{x}, \overleftrightarrow{y} \in P_{\overleftrightarrow{x},\overleftrightarrow{y}}$ for $\overleftrightarrow{x} \in \{\vec{s}, \overleftarrow{s}\}$ and $\overleftrightarrow{y} \in \{\vec{t}, \overleftarrow{t}\}$.*

Proof. It is easy to see that the existence of such four profiles implies that s and t are not weakly $\mathcal{P}$-nested.

So suppose for the other direction that s and t are not weakly $\mathcal{P}$-nested. Then, after potentially renaming s and t, there is a pair of profiles $P_{\vec{s},\vec{t}}, P_{\vec{s},\overleftarrow{t}}$ which is $< |s|$-indistinguishable such that $\vec{s}, \vec{t} \in P_{\vec{s},\vec{t}}$ and $\vec{s}, \overleftarrow{t} \in \mathcal{P}_{\vec{s},\overleftarrow{t}}$, and there is a pair of profiles $P_{\overleftarrow{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}}$ which is $< |s|$-indistinguishable such that $\overleftarrow{s}, \vec{t} \in P_{\overleftarrow{s},\vec{t}}$ and $\overleftarrow{s}, \overleftarrow{t} \in P_{\overleftarrow{s},\overleftarrow{t}}$.

If these profiles are not pairwise $< |s|$-indistinguishable, then there is a separation $r \in U$ such that $|r| < |s|$ and r distinguishes one of the profiles $P_{\vec{s},\vec{t}}, P_{\vec{s},\overleftarrow{t}}$ from one of the profiles $P_{\overleftarrow{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}}$. Since $P_{\vec{s},\vec{t}}, P_{\vec{s},\overleftarrow{t}}$ and $P_{\overleftarrow{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}}$ are $< |s|$-indistinguishable and $|r| < |s|$, this implies that there is an orientation $\overleftrightarrow{r}$ of r such that $\overleftrightarrow{r} \in P_{\vec{s},\vec{t}}, P_{\vec{s},\overleftarrow{t}}$ and $\overleftrightarrow{r}^* \in P_{\overleftarrow{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}}$.

Then however, we can consider the pair $\{\overleftrightarrow{r} \vee \vec{t}, \overleftrightarrow{r} \wedge \vec{t}\}$ of separations, and observe that one of these two separations has order $< |t| = |s|$ by submodularity. This separation of order less than $|t|$ then needs to distinguish $P_{\vec{s},\vec{t}}$ from $P_{\vec{s},\overleftarrow{t}}$ or $P_{\overleftarrow{s},\vec{t}}$ from $P_{\overleftarrow{s},\overleftarrow{t}}$ by the definition of a profile. This contradicts the fact that these two pairs are each $< |t|$-indistinguishable. $\square$

In order to be able to use Lemma 12 to prove Theorem 13, we need to show that the family of the sets $\mathcal{A}_{P,P'}$ of separations efficiently distinguishing the pair P, P' of profiles in $\mathcal{P}$ weakly splinters with respect to $\mathcal{P}$. For this we will need some analogue to the fish Lemma 2.3.1, i.e. we need to find a separation which can be used as a corner of two separation s and t which are not weakly $\mathcal{P}$-nested. We want to take as such a separation one of the two separation $\vec{s} \vee \vec{t}$ and $\vec{s} \wedge \vec{t}$. Thus, we want to show that this separation, at least in the cases where we need to find a corner of s and t, is $\mathcal{P}$-nested with all separations which are $\mathcal{P}$-nested with s and t. As our definition of when two separations r and s are $\mathcal{P}$-nested depends on the relation between the orders of r and s, we will, in order to show this, need to distinguish some cases. Our first lemma guarantees that, given two separations s and t such that $|s| \leqslant |t|$, a separation from $\{\vec{s} \vee \vec{t}, \vec{s} \wedge \vec{t}\}$ is still weakly $\mathcal{P}$-nested with all those separations of order higher than the order of t.

It will turn out that it is enough to show this in the case where the separation from $\{\vec{s} \vee \vec{t}, \vec{s} \wedge \vec{t}\}$ actually has the same order as t.

Lemma 4.2.8. *Let $r, s, t \in U$ such that $|s| \leqslant |t| \leqslant |r|$ and s and t are both weakly $\mathcal{P}$-nested with r. Let $\vec{c} \in \{\vec{s} \vee \vec{t}, \vec{s} \wedge \vec{t}\}$ such that $|c| = |t|$, then there is an orientation $\overleftrightarrow{c}$ of c such that any two profiles P, Q from $\mathcal{P}$ which are $< |c| = |t|$-indistinguishable and are distinguished by r do not both contain $\overleftrightarrow{c}$.*

Proof. There are orientations $\overleftrightarrow{s}$ and $\overleftrightarrow{t}$ of s and t witnessing that they are both weakly $\mathcal{P}$-nested with r. If c has an orientation $\overleftrightarrow{c}$ such that $\overleftrightarrow{c} \geqslant \overleftrightarrow{s}$ or $\overleftrightarrow{c} \geqslant \overleftrightarrow{t}$, say $\overleftrightarrow{c} \geqslant \overleftrightarrow{s}$, then clearly there cannot be two profiles P and Q in $\mathcal{P}$ which are $< |c|$-indistinguishable, distinguished by r and both contain $\overleftrightarrow{c}$: both P and Q would then by consistency also need to contain $\overleftrightarrow{s}$ which contradicts the fact that $\overleftrightarrow{s}$ witnesses that s and r are weakly $\mathcal{P}$-nested.

So we may suppose that c has an orientation such that $\overleftrightarrow{c} = \overleftrightarrow{s} \wedge \overleftrightarrow{t}$, in particular either $\overleftrightarrow{s} = \vec{s}$ and $\overleftrightarrow{t} = \vec{t}$, or $\overleftrightarrow{s} = \overleftarrow{s}$ and $\overleftrightarrow{t} = \overleftarrow{t}$. Let us suppose without loss of generality the first of the two, thus $\overleftrightarrow{c} = \vec{c} = \vec{s} \wedge \vec{t}$. We claim that $\vec{c}$ is the desired orientation of c.

Otherwise, there is a pair P, Q of profiles in $\mathcal{P}$ that is $< |c| = |t|$-indistinguishable, distinguished by r, and satisfies $\vec{c} \in P, Q$. Thus, by the profile property, $\vec{s} \in P$ or $\vec{t} \in P$, and $\vec{s} \in Q$ or $\vec{t} \in Q$. We may suppose that, after potentially changing the roles of P and Q, we have that $\vec{s} \in P$ and $\vec{t} \in Q$, but $\overleftarrow{s} \in P$ and $\overleftarrow{t} \in Q$: neither can both P and Q contain $\vec{s}$ as $\vec{s}$ witnesses that s and r are weakly $\mathcal{P}$-nested, nor can they both, by the same reason, contain $\vec{t}$.

This implies that $|s| = |t|$ since P and Q are $< |t|$ indistinguishable. Thus, we may suppose without loss of generality that $\vec{r} \in P$ and $\overleftarrow{r} \in Q$, as r distinguishes P and Q. Now since $|s| = |t|$, we may consider the profiles $P_{\vec{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}} \in \mathcal{P}$ from Lemma 4.2.7, and 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 u has an orientation $\overleftrightarrow{u}$ such that $\overleftrightarrow{u} \in P, Q$ and $\overleftrightarrow{u}^* \in P_{\vec{s},\vec{t}}, P_{\overleftarrow{s},\overleftarrow{t}}$. But then, by submodularity, one of the two separations $\vec{u} \vee \vec{s}$ and $\vec{u} \wedge \vec{s}$ has order less than $|s|$ and, by the definition of a profile, distinguishes P from Q or $P_{\vec{s},\vec{t}}$ from $P_{\overleftarrow{s},\overleftarrow{t}}$, which is a contradiction.

Thus, either $\vec{r} \in P_{\overleftarrow{s},\overleftarrow{t}}$ or $\overleftarrow{r} \in P_{\overleftarrow{s},\overleftarrow{t}}$. In the first case, $P_{\overleftarrow{s},\overleftarrow{t}}$ together with Q contradicts the assumption that $\vec{t}$ witnessed that t and r are weakly $\mathcal{P}$-nested, whereas $\overleftarrow{r} \in P_{\overleftarrow{s},\overleftarrow{t}}$ would result in $P_{\overleftarrow{s},\overleftarrow{t}}$ together with P contradicting the assumption that $\vec{s}$ witnessed that s and r are weakly $\mathcal{P}$-nested.

Thus, there cannot be such profiles P and Q and thus $\overleftrightarrow{c} = \vec{c}$ is as claimed. $\square$

The next three lemmas are needed to show a corresponding statement for separations r of order at most the order of t.

The first one will allow us to deduce in some cases that, if r is weakly $\mathcal{P}$-nested with both s and t, then we may assume that there is one orientation of r which witnesses that r is weakly $\mathcal{P}$-nested with both s and t.

Lemma 4.2.9. *Let $r, s, t \in U$ be separations such that $|r| \leqslant |s| \leqslant |t|$. Suppose that s and t are not weakly $\mathcal{P}$-nested, s and t are both weakly $\mathcal{P}$-nested with r, and there is a pair P, Q of profiles which is efficiently distinguished by t and not distinguished by s. Then there is an orientation $\overleftrightarrow{r}$ of r such that there is no pair of profiles P_1, P_2 from $\mathcal{P}$ which is $< |r|$-indistinguishable and distinguished by one of s and t such that both P_1 and P_2 contain $\overleftrightarrow{r}$.*

Proof. Since P and Q are not distinguished by s, they both contain the same orientation of s, say $\overleftrightarrow{s}$. Since s and t are not weakly $\mathcal{P}$-nested, there needs to be another pair P', Q' of $< |s|$-indistinguishable profiles which is distinguished by t, say because of $\vec{t} \in P'$ and $\overleftarrow{t} \in Q'$, and which both contain $\overleftrightarrow{s}^*$. Our general prove strategy will be to show that the four profiles P, Q, P', Q' are all pairwise $< |s|$-indistinguishable. This will then imply that they all four need to orient r, and since r is weakly $\mathcal{P}$-nested with s and t, we will then be able to analyse which orientation of r they need to contain, which will yield the claim.

So let us suppose without loss of generality that $\vec{t} \in P$ and $\overleftarrow{t} \in Q$.

We claim that P, Q, P', Q' are pairwise $< |s|$ indistinguishable. Indeed, if u is a separation of order less than $|s|$ which distinguish some of the two, then there is, as P, Q and P', Q' are pairwise $< |s|$-indistinguishable, an orientation $\overleftrightarrow{u}$ such that $\overleftrightarrow{u} \in P, Q$ and $\overleftrightarrow{u}^* \in P', Q'$. Thus, there is a corner of u and t which either contradicts the fact that P and Q are $< |t|$-indistinguishable or the fact that P', Q' are $< |s|$ indistinguishable.

Now all four profiles P, Q, P', Q' orient r. Thus, if $|r| < |s|$ we are already done: then they all need to contain the same orientation of r, say $\overleftrightarrow{r}^*$. Thus, $\overleftrightarrow{r}^*$ can neither witness that r and s are weakly $\mathcal{P}$-nested, nor that r and t are weakly $\mathcal{P}$-nested. Thus, $\overleftrightarrow{r}$ will need to witness both, which yields the claim.

So suppose in the following that $|r| = |s|$. Let us denote the orientation of r contained in P as $\overleftrightarrow{r}$. Now if $\overleftrightarrow{r} \in Q'$, then this would imply that $\overleftrightarrow{r}$ can neither witness that r and s are weakly $\mathcal{P}$-nested, nor that r and t are, as P and Q' are $< |r|$-indistinguishable and distinguished by s and t. Thus, $\overleftrightarrow{r}^*$ would need to witness both, and would thus be the desired orientation of r. So suppose that $\overleftrightarrow{r}^* \in Q'$.

Now either $\overleftrightarrow{r} \in P'$ or $\overleftrightarrow{r}^* \in P'$. In the first case, we observe that then $\overleftrightarrow{r}^*$ needs to be the orientation of r which witnesses that r and s are weakly $\mathcal{P}$-nested. Since Q and Q' are also distinguished by s, this then implies that also $\overleftrightarrow{r} \in Q$. Since Q and P' are distinguished by t, this then implies that $\overleftrightarrow{r}^*$ also witnesses that r and s are weakly $\mathcal{P}$-nested, and is thus again the required orientation of r.

In the second case, $\overleftrightarrow{r}^* \in P'$, we observe that, since t distinguishes P' and Q', the separation $\overleftrightarrow{r}$ needs to be the orientation of r which witnesses that r and t are weakly $\mathcal{P}$-nested. Thus, as before, since P and Q are also distinguished by t we observe that $\overleftrightarrow{r}^* \in Q$. Hence, as P' and Q are also distinguished by s we note that $\overleftrightarrow{r}$ needs to be the orientation of r which witnesses that r and s are weakly $\mathcal{P}$-nested. Therefore, $\overleftrightarrow{r}$ is the required orientation of r. □

We can now use this lemma to show that, given two separations s and t, in some cases a separation from $\{\vec{s} \vee \vec{t}, \vec{s} \wedge \vec{t}\}$ will by weakly $\mathcal{P}$-nested with those separations r which are weakly $\mathcal{P}$-nested with s and t and have an order between the order of s and t.

Lemma 4.2.10. *Let $s, r, t \in U$ be separations such that $|s| \leqslant |r| \leqslant |t|$, s and t are both weakly $\mathcal{P}$-nested with r, and t efficiently distinguishes a pair P, Q of profiles from $\mathcal{P}$ not distinguished by s. Let $\vec{c} \in \{\vec{s} \vee \vec{t}, \vec{s} \wedge \vec{t}\}$ such that $|\vec{c}| = |\vec{t}|$, then there is an orientation $\overleftrightarrow{r}$ of r such that any two profiles P, Q from $\mathcal{P}$ which are $< |\vec{r}|$-indistinguishable and are distinguished by c do not both contain $\overleftrightarrow{r}$.*

Proof. There is an orientation $\overleftrightarrow{r}$ of r witnessing that r and t are weakly $\mathcal{P}$-nested. We claim that this orientation is the required orientation of r. Otherwise, there is a pair P', Q' of $< |r|$-indistinguishable profiles that are distinguished by c and both contain $\overleftrightarrow{r}$. Since this pair is distinguished by c, but not distinguished by t, the profiles P' and Q' are, by the definition of a profile, distinguished by s, thus $|s| = |r|$.

Since $|s| = |r|$ and s and r are, by assumption, weakly $\mathcal{P}$-nested, we thus also find an orientation of r witnessing that r and s are weakly $\mathcal{P}$-nested. By Lemma 4.2.9, this orientation needs to be $\overleftrightarrow{r}$ as well. But this contradicts the choice of P', Q' since they both contain $\overleftrightarrow{r}$ and are distinguished by s. Thus, $\overleftrightarrow{r}$ is the required orientation of r. □

Finally, we can show that also those separations r of order less than the order of s and t that are weakly $\mathcal{P}$-nested with s and t will in some cases also be weakly $\mathcal{P}$-nested with the chosen separation from $\{\vec{s} \vee \vec{t}, \vec{s} \wedge \vec{t}\}$. This lemma actually follows directly from Lemma 4.2.9:

Lemma 4.2.11. *Let $s, r, t \in U$ be separations such that $|r| \leqslant |s| \leqslant |t|$, s and t are both weakly $\mathcal{P}$-nested with r, and t efficiently distinguishes a pair P, Q of profiles from $\mathcal{P}$ not distinguished by s. Let $\vec{c} \in \{\vec{s} \vee \vec{t}, \vec{s} \wedge \vec{t}\}$ such that $|\vec{c}| = |\vec{t}|$, then there is an orientation $\overleftrightarrow{r}$ of r such that any two profiles P, Q from $\mathcal{P}$ which are $< |\vec{r}|$-indistinguishable and are distinguished by c do not both contain $\overleftrightarrow{r}$.*

Proof. There are orientations of r which witness that r is weakly $\mathcal{P}$-nested with s and t. By Lemma 4.2.9, these orientations needs to coincide and equal $\overleftrightarrow{r}$, say. Thus, $\overleftrightarrow{r}$ is clearly the desired orientation of r, by the definition of a profile. □

We are now ready to prove the main Theorem 13 of this section:

Proof of Theorem 13. We consider, for $P, P' \in \mathcal{P}$, the set $\mathcal{A}_{P,P'}$ of those unoriented directed separations efficiently distinguishing P and P'. We want to show that the family of these $\mathcal{A}_{P,P'}$ splinters with respect to relation of being weakly $\mathcal{P}$-nested. In order to do so, let $s \in \mathcal{A}_{P,P'}, t \in \mathcal{A}_{Q,Q'}$, and suppose that $|s| \leqslant |t|$ and that s and t are not weakly $\mathcal{P}$-nested.

If s distinguishes Q and Q', then we can take $s \in \mathcal{A}_{Q,Q'}$ as the required corner of s and t. So we may suppose that there is an orientation $\overleftrightarrow{s}$ of s such that $\overleftrightarrow{s} \in Q, Q'$. Let us assume that $\overleftrightarrow{s} = \vec{s}$, the case $\overleftrightarrow{s} = \overleftarrow{s}$ is symmetric.

Let us now first consider the case that $|s| < |t|$. Since s and t are not weakly $\mathcal{P}$-nested, there needs to be a pair $Q_{\overleftarrow{s}}, Q'_{\overleftarrow{s}}$ of profiles which are $< |s|$-indistinguishable such that $\overleftarrow{s} \in Q_{\overleftarrow{s}}, Q'_{\vec{s}}$ and $\vec{t} \in Q_{\overleftarrow{s}}$, and $\overleftarrow{t} \in Q'_{\overleftarrow{s}}$. If $|\vec{s} \wedge \vec{t}| < |s|$, then this separation $\vec{s} \wedge \vec{t}$ would, by the profile property and consistency, also distinguish $Q_{\overleftarrow{s}}$ and $Q'_{\overleftarrow{s}}$ and would thus contradict the assumption that these two profiles are $< |s|$-indistinguishable. Thus, $|\vec{s} \wedge \vec{t}| \geqslant |s|$ and therefore, by submodularity, $|\vec{s} \vee \vec{t}| \leqslant |\vec{t}|$. Consequently, $\vec{s} \vee \vec{t}$ also distinguishes Q and Q' and thus $|\vec{s} \vee \vec{t}| = |t|$ as Q and Q' are $< |t|$-indistinguishable. Hence, $\vec{s} \vee \vec{t}$ is contained in $\mathcal{A}_{Q,Q'}$ and by the Lemmas 4.2.8, 4.2.10, and 4.2.11 the separation $\vec{s} \vee \vec{t}$ is weakly $\mathcal{P}$-nested with all separations r which are weakly $\mathcal{P}$-nested with both s and t. Thus, $\vec{s} \vee \vec{t}$ is the desired corner.

Now if $|s| = |t|$, there either are again profiles $Q_{\overleftrightarrow{s}}, Q'_{\overleftrightarrow{s}}$ as above, and we can perform the same argument, or we can perform the same argument after changing the roles of s and t.

Thus, by the splinter Lemma 12 there exists a weakly $\mathcal{P}$-nested set N of directed separations which meets all $\mathcal{A}_{P,P'}$, i.e. which efficiently distinguishes all the profiles from $\mathcal{P}$. □

So given a directed universe $\overleftrightarrow{U}$ with a submodular order function, we can distinguish all the profiles in $\mathcal{P}$ with such a weakly $\mathcal{P}$-nested set. However, what we actually want to have is a tree-like structure which distinguishes the profiles, like the tree-labellings for digraphs. So let us now see how we can construct such a structure from a weakly $\mathcal{P}$-nested set. The underlying principle of this structure will be the one of an S-tree, except that S now will be a set of unoriented directed separations.

Given a directed separation system $\overleftrightarrow{S}$, an *$\overleftrightarrow{S}$-tree*[3] consists of a tree T together with a map α from the set $\vec{E}(T)$ of the directed edges of T to the separations in $\overleftrightarrow{S}$ such that $\alpha(v,w) = \alpha(w,v)^*$ for every edge $(v,w) \in \vec{E}(T)$. Given a set $\mathcal{P}$ of distinguishable profiles in a directed universe $\overleftrightarrow{U}$, a $\overleftrightarrow{U}$-tree (T,α) is *for $\mathcal{P}$* if there exists a bijection $\beta : V(T) \to \mathcal{P}$ with the following property: given any two nodes v, w in $V(T)$ and the directed path vTw between v and w in T, we consider the set $\alpha(vTw) := \{\alpha(w_1, w_2) : (w_1, w_2) \in vTw\}$. We require that, given any separation $\overleftrightarrow{s} \in \alpha(vTw)$ of minimal possible order, this separation efficiently distinguishes $\beta(v)$ from $\beta(w)$, and it does so such that $\overleftrightarrow{s} \in \beta(w)$ and $\overleftrightarrow{s}^* \in \beta(v)$.

Let us now see how we can, given a directed universe $\overleftrightarrow{U}$, a set $\mathcal{P}$ of distinguishable profiles in $\overleftrightarrow{U}$ and a weakly $\mathcal{P}$-nested set of directed separations which efficiently distinguishes $\mathcal{P}$, construct a $\overleftrightarrow{U}$-tree for $\mathcal{P}$.

Theorem 4.2.12. *Let $\overleftrightarrow{U}$ be a directed universe with a submodular order function, $\mathcal{P}$ a set of pairwise distinguishable profiles in $\overleftrightarrow{U}$, and let N be a weakly $\mathcal{P}$-nested set of separations that efficiently distinguishes any two distinguishable profiles in $\mathcal{P}$. Then there is a $\overleftrightarrow{U}$-tree T for $\mathcal{P}$ such that the image of α is contained in $\overleftrightarrow{N}$.*

Proof. We may assume without loss of generality that the weakly $\mathcal{P}$-nested set N is chosen $\subseteq$-minimal with the property that it efficiently distinguishes any two distinguishable profiles in $\mathcal{P}$ efficiently.

We define our $\overleftrightarrow{U}$-tree T as follows: let $V(T) = \mathcal{P}$. Now for every $s \in N$, we add a single edge to T as follows. We pick an arbitrary pair of profiles P, Q in $\mathcal{P}$ which in N is efficiently distinguished by s only, add an edge between P and Q and define $\alpha(P,Q) = \overleftrightarrow{s}$, for the orientation $\overleftrightarrow{s}$ of s contained in Q. Consequently, we then set $\alpha(Q,P) = \overleftrightarrow{s}^*$, and note that $\overleftrightarrow{s}^* \in P$. Since N is chosen $\subseteq$-minimal, there need to exist, for every separation $s \in N$, a pair of profiles from $\mathcal{P}$ such

[3]We denote this object as $\overleftrightarrow{S}$-tree instead of S-tree, to highlight that we are working with a tree in the context of directed separation systems instead of an S-tree for an abstract separation system S.

that s is the only separation efficiently distinguishing that pair. Thus, in our construction we add exactly one edge for every separation from N.

First, we need to show that T is a tree. For this we define, given $k \in \mathbb{N}$ an equivalence relation $\sim_k$ on $\mathcal{P}$ by declaring that $P \sim_k Q$ if and only if P and Q induce the same profile of $\overleftrightarrow{S}_{k+1}$, i.e. the restriction $P\restriction_k$ of P to all separations of order at most k equals $Q\restriction_k$.

We now show inductively that every contraction $T_k = T/\sim_k$ is a tree, where this contraction, in principle, may contain multiple edges. Thus, $\vec{E}(T_k)$ contains precisely all those directed edges $\vec{e}$ from $\vec{E}(T)$ for which $\alpha(\vec{e})$ has order at most k.

For the induction start, observe that $\alpha(\vec{e})$ has order 0 for every directed edge $\vec{e}$ in T_0. Let us denote the image of $\vec{E}(T_0)$ under α as $\overleftrightarrow{N_0}$, this equals to the set of all separations of order 0 from $\overleftrightarrow{N}$. Each separation $\overleftrightarrow{s}$ in $\overleftrightarrow{N_0}$ defines a bipartition $(P^0_{\vec{s}}, P^0_{\overleftarrow{s}})$ of $\mathcal{P}_0 = \{P|_0 \,:\, P \in \mathcal{P}\}$ into the class $P^0_{\vec{s}} = \{P \in \mathcal{P}_0 \,:\, \vec{s} \in P\}$ of those profiles containing $\vec{s}$ and the class $P^0_{\overleftarrow{s}} = \{P \in \mathcal{P}_0 \,:\, \overleftarrow{s} \in P\}$ of those profiles containing $\overleftrightarrow{s}^*$.

We observe that the set N_0 is actually $\mathcal{P}_0$-nested, not just weakly $\mathcal{P}_0$-nested, i.e. that these bipartitions need to be nested in the usual way for a system of bipartitions: given $s, t \in N_0$ there are orientations $\overleftrightarrow{s}$ and $\overleftrightarrow{t}$ of s and t which witness that they are weakly $\mathcal{P}_0$-nested. This implies that there either is no profile in $\mathcal{P}_0$ contain $\overleftrightarrow{s}$ and $\overleftrightarrow{t}$, or no profile in $\mathcal{P}_0$ contain $\overleftrightarrow{s}$ and $\overleftrightarrow{t}^*$ or no profile in $\mathcal{P}_0$ containing $\overleftrightarrow{s}^*$ and $\overleftrightarrow{t}$ which shows that the bipartitions of $\mathcal{P}_0$ induced by s and t are nested in the usual way for bipartition.

Moreover, no two of the bipartitions induced by different $\overleftrightarrow{s}, \overleftrightarrow{t} \in N_0$ can coincide, since N was chosen $\subseteq$-minimal.

Let us denote as $\vec{\mathcal{B}}_0$ the set of all oriented bipartitions of $\mathcal{P}_0$ induced by one of the separations in $\overleftrightarrow{N_0}$, and note that $\vec{\mathcal{B}}_0$ equals a usual separation system of bipartitions of a set. Since every separation in $\overleftrightarrow{N_0}$ distinguishes some two profiles, all the bipartitions in $\vec{\mathcal{B}}_0$ are regular.

Every profile $P \in \mathcal{P}_0$ induces a consistent orientation of $\vec{\mathcal{B}}_0$ by orienting every bipartition in $\vec{\mathcal{B}}_0$ towards the side containing P. Note that these orientations are distinct for distinct profiles in $\mathcal{P}_0$. Let us denote as $\mathcal{O}_0$ the set of all the consistent orientations of $\vec{\mathcal{B}}_0$ induced by a profile in $\mathcal{P}_0$. As N was chosen $\subseteq$-minimal, the set $\vec{\mathcal{B}}_0$ is a $\subseteq$-minimal set of bipartitions which distinguishes all the orientations in $\mathcal{O}_0$. Thus, by Corollary 2.6.5, there is an order-respecting $\mathcal{B}$-tree (T', α') such that T' equals $\mathcal{O}_0$ and the image of α' corresponds exactly to $\vec{\mathcal{B}}_0$. As there is a natural bijection between $\mathcal{O}_0$ and $\mathcal{P}_0$, let us suppress this bijection and assume that the vertex set of T' equals exactly $\mathcal{P}_0$.

We claim that the edge set of T' equals the edge set of $T_0 = T/\sim_0$. Indeed, by the construction in Corollary 2.6.5, T' contains an edge between two profiles P and Q in $\mathcal{P}_0$ if and only if the bipartition $\alpha'(P, Q)$ is the only bipartition in $\vec{\mathcal{B}}_0$ distinguishing the orientations in $\mathcal{O}_0$ corresponding to P and Q. This is the case precisely if N_0 contains only one directed separation distinguishing P and Q. Conversely, for any two profiles which are not adjacent in T', there are at least two distinct bipartitions in $\vec{\mathcal{B}}_0$ which distinguish the orientations in $\mathcal{O}_0$ corresponding to P and Q and thus P and Q are also distinguished by at least two separations in N_0.

Now since T_0 contains an edge between two profiles only if these two are distinguished by only one separation from N_0, and conversely we contain one such edge for each separation from N_0, this implies that T_0 contains an edge between two profiles precisely if T' does. Moreover, the separations mapped to the orientated edges of T_0 (by α) and T' by α') coincide, i.e. for any $P, Q \in \mathcal{P}_0$ we have that the bipartition corresponding to $\alpha(P, Q)$ needs to equal $\alpha'(P, Q)$. In particular, T_0 is a tree as claimed.

For the induction step, let us denote, given a k-profile $P \in \mathcal{P}$ as $\mathcal{P}_P$ the set of $(k+1)$-profiles in $\mathcal{P}$ which induce P. It is enough to show that, for every such set $\mathcal{P}_P$, the subgraph induced by $\mathcal{P}_P$ on T_{k+1} is a tree, as T_{k+1} is obtained from T_k by expanding each node of T_k which corresponds to a k-profile P in $\mathcal{P}$ by that subgraph.

We first observe that for every separation s in N, there is at most one k-profile P in $\mathcal{P}$ such that s efficiently distinguishes profiles from $\mathcal{P}_P$: suppose s distinguishes some $(k+1)$-profiles $P_1, P_2 \in \mathcal{P}_P$ efficiently, and at the same time distinguish some $(k+1)$-profiles $Q_1, Q_2 \in \mathcal{P}_Q$ efficiently, for distinct profiles P and Q from $\mathcal{P}$. Since N distinguish all profiles from $\mathcal{P}$ from another, N contains a separation t that efficiently distinguishes P and Q. Moreover, as P and Q induce distinct profiles in $\mathcal{P}\restriction_k$, this separation t has order at most k. Thus, s and t cannot be weakly $\mathcal{P}$-nested by definition, which contradicts the fact that all separations from N are pairwise weakly $\mathcal{P}$-nested.

Thus, for $\mathcal{P}_k$ the set of k-profiles induced by a profile in $\mathcal{P}$, we can partition the set N_{k+1} of separations of order $k+1$ from N into the sets N_P of those separations efficiently distinguishing some two profiles form $\mathcal{P}_P$, i.e. $N_{k+1} = \bigcup_{P \in \mathcal{P}_k} N_P$ for $N_P = \{s \in N_{k+1} \,:\, s \text{ distinguishes two profiles } P', P'' \in \mathcal{P}_P\}$.

Now given any k-profile P from $\mathcal{P}$, the set N_P again induces a nested set of bipartitions of $\mathcal{P}_P$, and, exactly as in the induction start, we can show that the subgraph of T induced on $\mathcal{P}_P$ is a tree, again using the separation system of the induced bipartitions and Corollary 2.6.5. Thus, T is indeed a tree.

Let us now see that for any two nodes P, Q in $V(T)$, any separation $\overleftrightarrow{s} \in \alpha(PTQ)$ of minimal possible order efficiently distinguishes P from Q, and does so such that $\overleftrightarrow{s} \in Q$ and $\overleftrightarrow{s}^* \in P$. For this we note that either P and Q induce distinct profiles in $\mathcal{P}_0$, or there needs to be some $k \in \mathbb{N}$ such that P and Q induce the same profile P_k in $\mathcal{P}_k$. If they induce distinct profiles in $\mathcal{P}_0$, they correspond to distinct nodes in T_0, and any edge on the path between P and Q which corresponds to a separation of order 0 is also contained in T_0. Since the separations in T_0 correspond to the bipartitions in (T', α'), any separation along the directed path from P to Q distinguishes P and Q, and does so such that, for $(v, w) \in PTQ$, the separation $\alpha(v, w)$ is contained in Q. If on the other hand P and Q induced the same k-profile P_k, but distinct profiles of order $(k+1)$, then we can perform the exactly same argument as above inside the subtree of T_{k+1} induced on $\mathcal{P}_{P_k}$. □

Now Theorem 13 and Theorem 4.2.12 together imply Theorem 4.2.1.

Proof of Theorem 4.2.1. Since every tangle is a regular profile, Theorem 13 and Theorem 4.2.12 together imply that we find a $\overleftrightarrow{U}$-tree T for $\mathcal{T}$ so that the image of α is contained in $\overleftrightarrow{N}$. Let us denote the bijection between T and $\mathcal{T}$ given by that tree as β. Let us define a map $\gamma : E(T) \to S(G)$ by letting $\gamma(e)$ be the unoriented separation corresponding to $\{\alpha(\vec{e}), \alpha(\overleftarrow{e})\}$, for $\vec{e}, \overleftarrow{e}$ the two orientations of e. Then (T, β, γ) is the desired tree-labelling. □

4.3 A canonical tree-of-tangles theorem for structurally submodular separation systems

4.3.1 Introduction

In this section we present the following canonical version of Theorem 1.1.3:

Theorem 14. *Let $\vec{S}$ be a structurally submodular separation system and $\mathcal{P}$ a set of profiles of S. Then there is a nested set $N = N(\vec{S}, \mathcal{P}) \subseteq S$ which distinguishes $\mathcal{P}$. This $N(\vec{S}, \mathcal{P})$ can be chosen canonically: if $\varphi\colon \vec{S} \to \vec{S'}$ is an isomorphism of separation systems and $\mathcal{P}' := \{\varphi(P) : P \in \mathcal{P}\}$ then $\varphi(N(\vec{S}, \mathcal{P})) = N(\vec{S'}, \mathcal{P}')$.*

There are some technical subtleties in the formulation of Theorem 14 due to the fact that neither the profile property nor structural submodularity need be preserved by isomorphisms of separation systems. To avoid these difficulties, we obtain Theorem 14 by first establishing the following more general but somewhat more technical result, which slightly weakens the definitions of submodularity and profiles in order to make them compatible with such isomorphisms:

Theorem 4.3.1. *Let $\vec{S}$ be a separation system and $\mathcal{P}$ a collection of consistent orientations of $\vec{S}$ such that $\vec{S}$ is $\mathcal{P}$-submodular. Then there is a nested set $N = N(\vec{S}, \mathcal{P}) \subseteq S$ which distinguishes $\mathcal{P}$. This $N(\vec{S}, \mathcal{P})$ can be chosen canonically: if $\varphi\colon \vec{S} \to \vec{S'}$ is an isomorphism of separation systems and $\mathcal{P}' := \{\varphi(P) : P \in \mathcal{P}\}$, then $\varphi(N(\vec{S}, \mathcal{P})) = N(\vec{S'}, \mathcal{P}')$.*

Theorem 14 is then an immediate corollary of Theorem 4.3.1.

This section is structured as follows: in Section 4.3.2 we introduce the new definition required for Theorem 4.3.1 and show that Theorem 14 indeed is an immediate corollary of Theorem 4.3.1. In Section 4.3.3 we prove Theorem 14 and Theorem 4.3.1.

4.3.2 Submodularity with respect to a set of profiles

The first hurdle to overcome when aiming for a canonical version of Theorem 1.1.3 is to pin down what exactly 'canonical' ought to mean. At first glance this is obvious: the construction of the nested set N shall use only invariants of $\vec{S}$ and $\mathcal{P}$, that is, properties which are preserved by isomorphisms of separation systems. This approach, however, runs into a subtle difficulty: the definitions of both structural submodularity and profiles depend on $\vec{S}$ being embedded into an ambient universe of separations, whose existence Theorem 1.1.3 implicitly assumes. An isomorphism $\varphi\colon \vec{S} \to \vec{S'}$ of separation systems, though, need not preserve such an embedding, which leads to the undesirable situation that a construction isomorphic to that of N in $\vec{S}$ could be carried out in $\vec{S'}$, even though Theorem 14 may not be directly applicable to $\vec{S'}$ due to differences in their embeddings into ambient lattices.

To make our canonical version of Theorem 1.1.3 as widely applicable as possible, and to keep the definition of canonicity as straightforward and clean as possible, we must therefore tweak the assumptions of structural submodularity and profiles of S in such a way that they no longer depend on any embedding into a universe of separations, and are themselves invariants of isomorphisms between separation systems. This is made possible by the following observation: the proof of Theorem 1.1.3 makes use of the assumptions that $\vec{S}$ is submodular and $\mathcal{P}$ a set of profiles solely to deduce that whenever some $\vec{r}$ and $\vec{s}$ in $\vec{S}$ distinguish some two profiles in $\mathcal{P}$, then either their meet or their join (as provided by the ambient lattice) is contained in $\vec{S}$ and likewise distinguishes that pair of profiles.

For our canonical Theorem 4.3.1 we will thus eliminate the need for an ambient universe by asking of $\vec{S}$ and $\mathcal{P}$ that they have this property, with the meets or joins now being taken directly in the poset $\vec{S}$. Expressed solely in terms of $\vec{S}$ and $\mathcal{P}$, this 'richness' property is then preserved by isomorphisms of separation systems, independently of any embeddings into lattice structures. This solves the minor problem in the formulation of Theorem 14 of $\vec{S'}$ not meeting the assumption of the theorem despite being isomorphic to a separation system which does. Theorem 14 is then obtained as a corollary of Theorem 4.3.1.

Let $\vec{S}$ be a separation system and $\mathcal{P}$ a set of consistent orientations of $\vec{S}$. Recall that given a set $M \subseteq \vec{S}$ of oriented separations, an element $\vec{r} \in \vec{S}$ is an *infimum* of M in $\vec{S}$ if $\vec{r} \leqslant \vec{s}$ for each $\vec{s} \in M$ and additionally $\vec{r} \geqslant \vec{t}$ whenever $\vec{t} \in \vec{S}$ is such that $\vec{t} \leqslant \vec{s}$ for all $\vec{s} \in M$. Dually, an element $\vec{r} \in \vec{S}$ is a *supremum* of M in $\vec{S}$ if $\vec{r} \geqslant \vec{s}$ for each $\vec{s} \in M$ and additionally $\vec{r} \leqslant \vec{t}$ whenever $\vec{t} \in \vec{S}$ is such that $\vec{t} \geqslant \vec{s}$ for all $\vec{s} \in M$. In general a set $M \subseteq \vec{S}$ need not have such an infimum or supremum in $\vec{S}$.

Given two separations $\vec{r}$ and $\vec{s}$ in $\vec{S}$ we denote the infimum and supremum of $\{\vec{r}, \vec{s}\}$ in $\vec{S}$ by $\vec{r} \wedge \vec{s}$ and $\vec{r} \vee \vec{s}$, respectively, if those exist. Observe that $(\vec{r} \vee \vec{s})^* = \overleftarrow{r} \wedge \overleftarrow{s}$.

If $\vec{r}$ and $\vec{s}$ have a supremum $\vec{r} \vee \vec{s}$ in $\vec{S}$, and every $P \in \mathcal{P}$ containing both $\vec{r}$ and $\vec{s}$ also contains $\vec{r} \vee \vec{s}$, then we call $\vec{r} \vee \vec{s}$ a *$\mathcal{P}$-join* of $\vec{r}$ and $\vec{s}$ in $\vec{S}$. Dually we call $\vec{r} \wedge \vec{s}$ a *$\mathcal{P}$-meet* of $\vec{r}$ and $\vec{s}$ if $(\vec{r} \wedge \vec{s})^* \in P$ for each $P \in \mathcal{P}$ containing both $\overleftarrow{r}$ and $\overleftarrow{s}$.

Finally, we say that $\vec{S}$ is *$\mathcal{P}$-submodular* if every two crossing separations $\vec{r}$ and $\vec{s}$ in $\vec{S}$ have a $\mathcal{P}$-join or $\mathcal{P}$-meet in $\vec{S}$. For the remainder of this section we assume $\vec{S}$ to be $\mathcal{P}$-submodular.

Observe that $\mathcal{P}$-submodularity is preserved by isomorphisms of separation systems: if $\varphi\colon \vec{S} \to \vec{S'}$ is an isomorphism and $\mathcal{P}' := \{\varphi(P) \,:\, P \in \mathcal{P}\}$, then $\vec{S'}$ is $\mathcal{P}'$-submodular. Using this notion of submodularity we can therefore meaningfully express canonicity in the context of Theorem 1.1.3.

Note, however, that we are not assuming the elements of $\mathcal{P}$ to be profiles of $\vec{S}$: this is precisely because we prove Theorem 4.3.1 without an ambient lattice structure, which would be necessary to define profiles. Therefore, Theorem 4.3.1 improves on Theorem 1.1.3 not only by offering a canonical way of constructing N, but also by being applicable to an even larger number of separation systems.

Before getting to the proof of Theorem 4.3.1 itself, let us demonstrate that it is in fact a strengthening of Theorem 1.1.3 by showing that Theorem 4.3.1 implies Theorem 14. The following lemma does just that by proving that in the setting of Theorem 1.1.3 the assumptions of Theorem 4.3.1 are satisfied:

Lemma 4.3.2. *Let $\vec{S}$ be a structurally submodular separation system inside some universe $\vec{U}$ of separations and $\mathcal{P}$ a set of profiles of $\vec{S}$. Then $\vec{S}$ is $\mathcal{P}$-submodular.*

Proof. We must show that any $\vec{r}$ and $\vec{s}$ in $\vec{S}$ have a $\mathcal{P}$-meet or $\mathcal{P}$-join. So let $\vec{r}$ and $\vec{s}$ in $\vec{S}$ be given. Since $\vec{S}$ is structurally submodular, it contains the infimum $\vec{r} \wedge \vec{s}$ or the supremum $\vec{r} \vee \vec{s}$ of $\vec{r}$ and $\vec{s}$ in $\vec{U}$. Let us assume the latter; the other case is dual. Then $\vec{r} \vee \vec{s}$ is also the supremum of $\vec{r}$ and $\vec{s}$ as taken in $\vec{S}$.

Moreover, by the profile property, every $P \in \mathcal{P}$ containing both $\vec{r}$ and $\vec{s}$ also contains $\vec{r} \vee \vec{s}$, making this a $\mathcal{P}$-join of $\vec{r}$ and $\vec{s}$ in $\vec{S}$. □

4.3.3 Proof of Theorem 14 and Theorem 4.3.1

In this section we will prove Theorem 14 and Theorem 4.3.1. To this end, for the remainder of this section let $\vec{S}$ be a separation system and let $\mathcal{P}$ be a set of consistent orientations of $\vec{S}$ such that $\vec{S}$ is $\mathcal{P}$-submodular.

As usual, an important tool in our proof of Theorem 4.3.1 is the fish Lemma 2.3.1. Since Lemma 2.3.1 is formulated in the context of a separation system that is contained in a universe of separations, we need to prove our own version of this lemma:

Lemma 4.3.3 (See also Lemma 2.3.1,[21, Lemma 3.2])**.** *Let $r, s \in S$ be two crossing separations in S and let $t \in S$ be a separation that is nested with both r and s. Given orientations $\vec{r}$ and $\vec{s}$ of r and s such that there exists a supremum $\vec{r} \vee \vec{s}$ of $\vec{r}$ and $\vec{s}$ in $\vec{S}$, the separation $\vec{r} \vee \vec{s}$ is nested with t. The same is true for $\vec{r} \wedge \vec{s}$.*

Proof. If t has an orientation $\vec{t}$ such that $\vec{t} \leqslant \vec{r}$ or $\vec{t} \leqslant \vec{s}$, then clearly $\vec{t} \leqslant (\vec{r} \vee \vec{s})$. Otherwise, since r and s cross, there must be an orientation $\vec{t}$ of t such that $\vec{r} \leqslant \vec{t}$ and $\vec{s} \leqslant \vec{t}$. Thus, by the fact that $\vec{r} \vee \vec{s}$ is the supremum of $\vec{r}$ and $\vec{s}$ in $\vec{S}$, we have that $(\vec{r} \vee \vec{s}) \leqslant \vec{t}$. □

Let us say that a separation $\vec{s} \in \vec{S}$ is *exclusive (for $\mathcal{P}$)* if it lies in exactly one orientation in $\mathcal{P}$. If $P \in \mathcal{P}$ is the orientation containing an exclusive separation $\vec{s}$, then we might also say that $\vec{s}$ is *P-exclusive (for $\mathcal{P}$)*. Observe that if $\vec{r}$ is P-exclusive for $\mathcal{P}$, then so is every $\vec{s} \in P$ with $\vec{r} \leqslant \vec{s}$.

For each $P \in \mathcal{P}$ let M_P consist of the maximal elements of the set of all P-exclusive separations. Equivalently, M_P is the set of all maximal elements of P that are exclusive for $\mathcal{P}$.

Our strategy for proving Theorem 4.3.1 will be to canonically pick nested P-exclusive representatives of all orientations $P \in \mathcal{P}$ that contain exclusive separations, then discard from $\mathcal{P}$ and $\vec{S}$ all those orientations P for whom we selected a representative and all those separations not nested with these representatives, respectively. Iterating this procedure will yield the canonical nested set.

In order for this strategy to work we must ensure that the sets M_P are not all empty. Our first lemma addresses this:

Lemma 4.3.4. *If $\mathcal{P}$ and S are non-empty, then some M_P is non-empty.*

In the case of $\vec{S}$ being submodular in some universe of separations $\vec{U}$ and $\mathcal{P}$ being a set of profiles of $\vec{S}$, the existence of exclusive separations and thus Lemma 4.3.4 is actually an immediate consequence of Theorem 1.1.3: if $N \subseteq S$ is a nested set which distinguishes $\mathcal{P}$, and each element of N distinguishes some pair of profiles in $\mathcal{P}$, then any maximal element of $\vec{N}$ is exclusive for $\mathcal{P}$. In other words, the separations labelling the incoming edges of leaves of the tree associated with N are exclusive. (See [22] for the precise relationship between nested sets and trees.)

However, since we are working with the more general notion of $\vec{S}$ being $\mathcal{P}$-submodular, we give an independent proof of Lemma 4.3.4.

Proof of Lemma 4.3.4. If $\mathcal{P}$ consists of only one orientation the assertion is trivial since S is non-empty. For $|\mathcal{P}| \geqslant 2$ we show the following stronger claim by induction on $|\mathcal{P}|$:

If $|\mathcal{P}| \geqslant 2$ there is for each $P \in \mathcal{P}$ a separation that is exclusive but not P-exclusive for $\mathcal{P}$.

For the base case $|\mathcal{P}| = 2$ observe that any separation distinguishing the two orientations in $\mathcal{P}$ has two exclusive orientations, one in each element of $\mathcal{P}$.

Suppose now that $|\mathcal{P}| > 2$ and that the claim holds for all non-singleton proper subsets of $\mathcal{P}$. Let $P \in \mathcal{P}$ be the given fixed orientation and set $\mathcal{P}' := \mathcal{P} \setminus \{P\}$. By the induction hypothesis applied to $\mathcal{P}'$ and an arbitrary orientation in $\mathcal{P}'$ there is an exclusive separation $\vec{r}$ for $\mathcal{P}'$, contained in some $Q \in \mathcal{P}'$. Applying the induction hypothesis again to $\mathcal{P}'$ and Q yields another separation $\vec{s}$ that is exclusive for $\mathcal{P}'$ and lies in some $Q' \in \mathcal{P}'$ with $Q \neq Q'$.

If $\vec{r}$ or $\vec{s}$ is also exclusive for $\mathcal{P}$, then we are done. So suppose not, that is, suppose we have $\vec{r}, \vec{s} \in P$. Then $r \neq s$, and hence $\vec{r}$ and $\vec{s}$ must be incomparable by the consistency of Q and Q'. If $\vec{r} \leqslant \overleftarrow{s}$, then $\overleftarrow{s}$ is Q-exclusive for $\mathcal{P}$. Moreover, $\overleftarrow{s} \leqslant \vec{r}$ is not possible by the consistency of P. Thus, we may assume that r and s cross.

By $\mathcal{P}$-submodularity of $\vec{S}$, there is a $\mathcal{P}$-join or a $\mathcal{P}$-meet of $\vec{r}$ and $\overleftarrow{s}$ in $\vec{S}$; by symmetry we may assume that there is a $\mathcal{P}$-join $(\vec{r} \vee \overleftarrow{s}) \in \vec{S}$. Since $\vec{s}$ is Q'-exclusive, we have $\overleftarrow{s} \in Q$ and hence $(\vec{r} \vee \overleftarrow{s}) \in Q$ by the fact that $\vec{r} \vee \overleftarrow{s}$ is the $\mathcal{P}$-join of $\vec{r}$ and $\overleftarrow{s}$. From $(\vec{r} \vee \overleftarrow{s}) \geqslant \vec{r}$ we infer that $\vec{r} \vee \overleftarrow{s}$ is Q-exclusive for $\mathcal{P}'$. Moreover, we cannot have $(\vec{r} \vee \overleftarrow{s}) \in P$: it would be inconsistent with $\vec{s} \in P$ as r and s cross.

Therefore, $\vec{r} \vee \overleftarrow{s}$ is exclusive but not P-exclusive for $\mathcal{P}$. □

We remark that, in the case of a submodular separation system $\vec{S}$ inside a universe of separations and a set $\mathcal{P}$ of profiles, the stronger assertion used for the induction hypothesis in this proof, too, can be established immediately using Theorem 1.1.3: for $|\mathcal{P}| \geqslant 2$ the tree associated with the nested set $N \subseteq S$ distinguishing $\mathcal{P}$ has at least two leaves, and hence some leaf for which the separation labelling its incoming edge does not lie in the fixed profile P.

Returning to the proof of Theorem 4.3.1, let us find a way to canonically pick representatives of those $P \in \mathcal{P}$ with non-empty M_P in such a way that these representatives are nested with each other. For the 'canonically'-part of this we will make use of the fact that the sets M_P themselves are invariants of $\mathcal{P}$ and $\vec{S}$. For the nestedness we start by showing that separations from different M_P's cannot cross at all:

Lemma 4.3.5. *For $P \neq P'$ all $\vec{r} \in M_P$ and $\vec{s} \in M_{P'}$ are pairwise nested.*

Proof. Suppose some $\vec{r} \in M_P$ and $\vec{s} \in M_{P'}$ cross. By $\mathcal{P}$-submodularity of $\vec{S}$, there is a $\mathcal{P}$-join $\vec{r} \vee \overleftarrow{s}$ or a $\mathcal{P}$-meet $\vec{r} \wedge \overleftarrow{s}$ in $\vec{S}$; by symmetry we may suppose that $(\vec{r} \vee \overleftarrow{s}) \in \vec{S}$. Then P, too, contains this separation since $\overleftarrow{s} \in P$. But $\vec{r} \vee \overleftarrow{s}$ is also P-exclusive and strictly larger than $\vec{r}$, a contradiction. □

It is possible, however, that the set M_P itself is not nested. In fact the elements of M_P all cross each other, unless $\mathcal{P} = \{P\}$: any $\vec{r}$ and $\vec{s}$ in M_P that are nested must point towards each other by maximality. But every other orientation in $\mathcal{P}$ contains both $\overleftarrow{r}$ and $\overleftarrow{s}$ and would then be inconsistent. If we want to represent a $P \in \mathcal{P}$ with non-empty M_P by an element of M_P, we are therefore limited to picking at most one element of M_P. However, there is no canonical way of singling out an element of M_P to be the representative of P; we must therefore find another way of choosing an invariant P-exclusive separation, using M_P only as a starting point.

For this we will show that each M_P has an infimum in $\vec{S}$ and that this infimum is again P-exclusive:

Lemma 4.3.6. *Let $P \in \mathcal{P}$ with $M_P \neq \emptyset$ and $\mathcal{P} \neq \{P\}$ be given. Then M_P has an infimum $\vec{s}_P$ in the poset $\vec{S}$, and $\vec{s}_P$ is P-exclusive for $\mathcal{P}$. Moreover, if some $t \in S$ is nested with M_P, then t is also nested with s_P.*

Proof. Fix an enumeration $M_P = \{\vec{r}_1, \ldots, \vec{r}_n\}$ and some $t \in S$ that is nested with M_P. We will show by induction on i that there is an infimum $\vec{s}_i \coloneqq \inf\{\vec{r}_1, \ldots, \vec{r}_i\}$ in $\vec{S}$; that this infimum is P-exclusive for $\mathcal{P}$; and that it is nested with t. This then yields the claim for $i = n$.

The case $i = 1$ is trivially true, so suppose that $i > 1$ and that we already know that the separation $\vec{s}_{i-1} = \inf\{\vec{r}_1, \ldots, \vec{r}_{i-1}\}$ is the infimum of $\vec{r}_1, \ldots, \vec{r}_{i-1}$ in $\vec{S}$, that it is P-exclusive, and that it is nested with t.

In the case that $s_{i-1} = r_i$ we have either $\vec{s}_{i-1} = \vec{r}_i$ or $\vec{s}_{i-1} = \overleftarrow{r}_i$. The latter of these is impossible, since P contains both of the P-exclusive separations $\vec{s}_{i-1}$ and $\vec{r}_i$. The former, however, gives that $\vec{r}_i$ is the infimum of $\vec{r}_1, \ldots, \vec{r}_i$ and thus as claimed by $\vec{r}_i \in M_P$.

So suppose that $s_{i-1} \neq r_i$. Let us first treat the case that $\vec{r}_i$ and $\vec{s}_{i-1}$ are nested. Clearly the two cannot point away from each other since P is consistent. If $\vec{r}_i$ and $\vec{s}_{i-1}$ are comparable, then one of the two is the infimum of $\vec{r}_i$ and $\vec{s}_{i-1}$ and thus the infimum of $\vec{r}_1, \ldots, \vec{r}_i$ in $\vec{S}$. Since both $\vec{s}_{i-1}$ and $\vec{r}_i$ are P-exclusive and nested with t, this infimum is thus as claimed. Finally, if $\vec{r}_i$ and $\vec{s}_{i-1}$ point towards each other, we obtain a contradiction: for then their inverses point away from each other, making every orientation in $\mathcal{P}$ other than P inconsistent. Thus, if $\vec{r}_i$ and $\vec{s}_{i-1}$ are nested the induction hypothesis holds for $\vec{s}_i$.

Let us now consider the case that $\vec{r}_i$ and $\vec{s}_{i-1}$ cross. Then there needs to be a $\mathcal{P}$-join or a $\mathcal{P}$-meet of $\vec{r}_i$ and $\vec{s}_{i-1}$. However, we cannot have a $\mathcal{P}$-join $\vec{r}_i \vee \vec{s}_{i-1}$ in $\vec{S}$ since this join would be P-exclusive and strictly larger than $\vec{r}_i \in M_P$. Therefore, there is a $\mathcal{P}$-meet $(\vec{r} \wedge \vec{s}_{i-1}) \in \vec{S}$. By consistency we have that $\vec{s}_i \in P$. Every orientation in $\mathcal{P}$ other than P contains $\overleftarrow{r}_i$ as well as $\overleftarrow{s}_{i-1}$ and hence $\overleftarrow{s}_i$ by the definition of $\mathcal{P}$-meet, which shows that $\vec{s}_i$ is P-exclusive. Finally, by Lemma 4.3.3, $\vec{s}_i$ is also nested with t. □

It remains to show that after picking as a representative for each $P \in \mathcal{P}$ with exclusive separations the infimum of M_P, the set of separations in $\vec{S}$ that are nested with all these representatives is still rich enough to distinguish all orientations in $\mathcal{P}$ for which we have not yet picked a representative.

For this let $\vec{S'} \subseteq \vec{S}$ be the system of all those separations that are nested with all M_P, and let $\mathcal{P}' \subseteq \mathcal{P}$ be the set of those orientations Q that have empty M_Q. Our next lemma says that if we restrict ourselves to $\vec{S'}$, we can still distinguish $\mathcal{P}'$:

Lemma 4.3.7. *The separation system $\vec{S'}$ is $\mathcal{P}'$-submodular and distinguishes $\mathcal{P}'$.*

Proof. The fact that $\vec{S'}$ is $\mathcal{P}'$-submodular is a direct consequence of Lemma 4.3.3: it implies that for $\vec{r}$ and $\vec{s}$ in $\vec{S'}$ any $\mathcal{P}$-meet or $\mathcal{P}$-join of them in $\vec{S}$ is contained in $\vec{S'}$. Since $\mathcal{P}' \subseteq \mathcal{P}$, this $\mathcal{P}$-join or $\mathcal{P}$-meet is in fact a $\mathcal{P}'$-join or $\mathcal{P}'$-meet of $\vec{r}$ and $\vec{s}$ in $\vec{S'}$.

To see that $\vec{S'}$ distinguishes $\mathcal{P}'$, let Q and Q' be distinct orientations in $\mathcal{P}'$; we shall show that some $\vec{s'} \in \vec{S'}$ distinguishes them. For this choose a separation $s \in S$ which distinguishes Q and Q' and which is nested with M_P for as many $P \in \mathcal{P}$ as possible. If s is nested with all M_P we are done; otherwise there is some $P \in \mathcal{P}$ for which s crosses some separation in M_P.

So suppose that there is a $P \in \mathcal{P}$ for which s is not nested with M_P. Among all $\vec{s'} \in \vec{S}$ which distinguish Q and Q' and which are nested with each $M_{P'}$ with which s is nested, pick a minimal $\vec{s'}$ with $\vec{s'} \in P$. We claim that this $\vec{s'}$ is nested with M_P, contradicting the choice of s.

To see this, suppose that $\vec{s'}$ crosses some $\vec{r} \in M_P$. Then there cannot exist a $\mathcal{P}$-join of $\vec{r}$ and $\vec{s'}$ in $\vec{S}$ since this join would be a strictly larger P-exclusive separation than $\vec{r}$. Hence, there is a $\mathcal{P}$-meet $\vec{r} \wedge \vec{s'}$ of $\vec{r}$ and $\vec{s'}$ in $\vec{S}$. By $P \notin \{Q, Q'\}$ we have that both Q and Q' contain $\overleftarrow{r}$, and hence this separation $\vec{r} \wedge \vec{s'}$ distinguishes Q and Q' as well: one of the two orientations contains $\vec{s'}$ and thus also $\vec{r} \wedge \vec{s'}$ by consistency. The other contains both $\overleftarrow{s'}$ and $\overleftarrow{r}$ and thus also $(\overleftarrow{r} \vee \overleftarrow{s'}) = (\vec{r} \wedge \vec{s'})^*$ by the fact that $\vec{r} \wedge \vec{s'}$ is the $\mathcal{P}$-meet of $\vec{r}$ and $\vec{s'}$. However, by Lemma 4.3.3 and Lemma 4.3.5, this $\vec{r} \wedge \vec{s'}$ would be nested with each $M_{P'}$ with which s was nested, while being strictly smaller than $\vec{s'}$, a contradiction. □

If M_P is non-empty let us write $\vec{s}_P$ for its infimum in $\vec{S}$ as in Lemma 4.3.6. We are now ready to prove Theorem 4.3.1 by induction.

Proof of Theorem 4.3.1. We proceed by induction on $|\mathcal{P}|$. If $|\mathcal{P}| \leqslant 1$ there is nothing to show, so suppose that $|\mathcal{P}| > 1$ and that the assertion holds for all proper subsets of $\mathcal{P}$.

Recall that $\vec{S'} \subseteq \vec{S}$ consists of all separations in $\vec{S}$ that are nested with all sets M_Q and that $\mathcal{P}' \subseteq \mathcal{P}$ is the set of all $Q \in \mathcal{P}$ with empty M_Q. Clearly both $\vec{S'}$ and $\mathcal{P}'$ are invariants of $\vec{S}$ and $\mathcal{P}$ since the sets M_Q themselves are invariants. For each non-empty M_Q let $\vec{s}_Q$ be its infimum in $\vec{S}$ as described in Lemma 4.3.6. Then

$$N_1 \coloneqq \{s_Q \;:\; Q \in \mathcal{P} \smallsetminus \mathcal{P}'\}$$

is clearly a canonical set. From Lemma 4.3.6 we further know that N_1 distinguishes all orientations in $\mathcal{P} \smallsetminus \mathcal{P}'$ from each other and from each orientation in $\mathcal{P}'$.

By Lemma 4.3.5, every element of M_P is nested with every element of $M_{P'}$ for all $P \neq P'$. Applying the 'moreover'-part of Lemma 4.3.6 twice thus implies that s_P is nested with every element of $M_{P'}$ and subsequently with $s_{P'}$. Therefore, N_1 is a nested set. Likewise, every separation in $\vec{S'}$ is nested with N_1.

Let us apply the induction hypothesis to $\mathcal{P}'$ in $\vec{S'}$, as made possible by the Lemmas 4.3.4 and 4.3.7, yielding a canonical nested set $N_2 \subseteq S'$ which distinguishes $\mathcal{P}'$. Since $\vec{S'}$ and $\mathcal{P}'$ themselves are invariants of $\vec{S}$ and $\mathcal{P}$, we have that the union $N_1 \cup N_2$ is the desired canonical nested set. □

Proof of Theorem 14. By Lemma 4.3.2, given a structurally submodular separation system $\vec{S}$ and a set $\mathcal{P}$ of profiles of S, we know that $\vec{S}$ is $\mathcal{P}$-submodular. Thus, Theorem 14 follows from Theorem 4.3.1. □

Acknowledgement We would like to thank one of the anonymous reviewers of the paper [36] on which this section of this thesis is based, for suggesting the concept of $\mathcal{P}$-submodularity, which led to Theorem 4.3.1.

4.4 Trees of tangles in infinite separation systems

4.4.1 Introduction

In Section 4.1 we introduced the 'splinter lemma', a unified lemma which implies the known tree-of-tangles theorems for finite separations systems. The merit of this lemma lies in the fact that, while it is strong enough to imply all these results, the proof of the lemma is simple, and its assumptions are easy to check.

Lemma 10 (Splinter Lemma)**.** *Let U be a universe of separations and let $\mathfrak{A} = (\mathcal{A}_i)_{i \leqslant n}$ be 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.*

Lemma 10, in a sense, is yet another step in a series of abstractions in the theory of tangles: rather than working with the tangles themselves, it operates just on the collection of sets of separations distinguishing a given pair of these.

Recall that Lemma 10 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 of this thesis we overcome these difficulties and present a way to obtain a version of Lemma 10 for infinite families of sets of separations which is as abstract and therefore as widely applicable, as our original Lemma 10 and which allow us to deduce the existing tree-of-tangles theorems in infinite graphs (see [9, 14]). It asks more of the sets $\mathcal{A}_i$ than our finite splinter lemma; in return however the set of separations we obtain will be canonical, i.e. invariant under isomorphisms:

Lemma 15. *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 $N((\, \varphi(\mathcal{A}_i) \,:\, i \in I\,)) = \varphi(N((\, \mathcal{A}_i \,:\, i \in I\,)))$.*

After stating some additional terminology as well as some basic facts in Section 4.4.2, we prove this statement in Section 4.4.3. Like Lemma 10, the statement of this lemma is a bit technical, as we want it to be as widely applicable as possible. We then show the usefulness of this abstract lemma throughout Section 4.4.4, by deducing the existing theorems about distinguishing tangles in infinite graphs from it.

As a simple example, we start with applying it to tangles in locally finite graphs. This application is straightforward and demonstrates a prototypical application of Lemma 15.

After that, we show that it is also possible to apply Lemma 15 to arbitrary infinite graphs. This application uses another new, and interesting, shift of perspective: as in Section 4.2, we cannot apply Lemma 15 directly to the sets of separations efficiently distinguishing two profiles since, in general, these do not splinter thinly. Instead, we need to consider a slightly different set, namely, the set of separators, to which Lemma 15 does apply:

Theorem 16. *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 $\mathcal{P}$.*

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 4.4.1 ([9, 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 [19, 59]. We deduce Theorem 4.4.1 from Theorem 16 in Section 4.4.4.

On the other hand, if we want to keep canonicity, we can use Theorem 16 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 4.4.2 ([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:*

(a) the tree of tree-decompositions distinguishes $\mathcal{P}$ efficiently;

(b) if $t \in V(T)$ has level k, then $(T_t, \mathcal{V}_t)$ contains only separations of order k;

(c) 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 deduce Theorem 4.4.2 from Theorem 16 in Section 4.4.4.

Theorem 16 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 4.4.1 and the canonical tree of tree-decompositions in Theorem 4.4.2.

Moreover, proving Theorem 4.4.1 or Theorem 4.4.2 by first proving Theorem 16 and then deducing them breaks up the proof nicely and is, in total, shorter than the original proofs from [9, 14].

4.4.2 Additional terminology and basic facts

Recall that a profile P in a graph G is *regular* if it does not contain any co-small 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-cutvertex. Formally, we can summarize this statement from [23] as follows:

Lemma 4.4.3 ([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 this section.

Lemma 4.4.4 ([56, 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 from an arbitrary separation system that are nested with a corner separation:

Lemma 4.4.5. *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 $\vec{t} \leqslant \overleftarrow{r} \leqslant (\vec{r} \wedge \vec{s})^*$. □

Moreover, we will use that the separations in a universe of separations with an order function that distinguish a given pair of profiles exhibit a lattice-like structure:

Lemma 4.4.6. *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 (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. □

Furthermore, we shall need a way to transition between separations and tree-decompositions in graphs. Such a method in finite graphs if for example provided in [13]. As it turns out, the ingredients of that proof together with Theorem 2.6.1 are all that is needed to show an analogous result for infinite graphs, which we shall present here.

We shall need the following lemma, whose proof is inspired by [13].

Lemma 4.4.7. *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 any ω-chain $(A_1, B_1) < (A_2, B_2) < \dots$ 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 from Theorem 2.6.1. Note that by [52, Theorem 3.9(iii)] any isomorphism between the edge tree sets of two distinct trees induces an isomorphism of the underlying trees.

Let α be the isomorphism from the edge tree set of T to N. Recall that, given some node $t \in T$ we denote as $\alpha(t)$ the set

$$\alpha(t) := \{\alpha(s,t) \,:\, (s,t) \in \vec{E}\}$$

of oriented separations. We define the bags of our tree-decomposition as $V_t := \bigcap_{(A,B)\in\alpha(t)} B$. 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 $v \in B \smallsetminus A$ 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 that ω-chain in O. Thus, (T1) holds.

The proof that (T2) holds 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 check that (T3) holds, 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 $t_2 \in P$. □

We remark that 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:

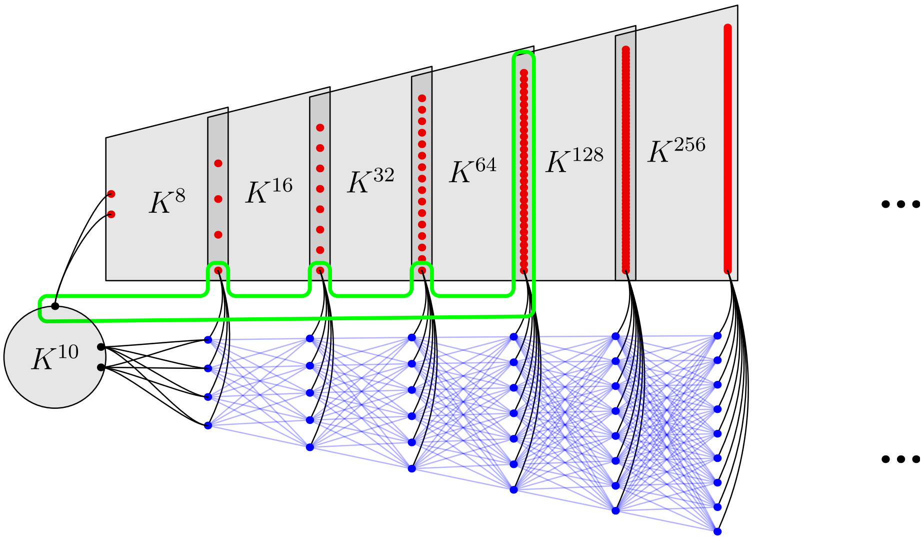

Figure 4.4: A locally finite graph where no tree-decomposition distinguishes all the robust regular bounded profiles efficiently. The green separator is the one of the only separation which efficiently distinguishes the profile induced by the K^{64} from the profile induced by the K^{128}.

Example 4.4.8. Consider the graph displayed in Fig. 4.4. 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.

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 $s_1, s_2, \ldots$, 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 4.4.7.

4.4.3 The thin splinter lemma

In this section we generalize the finite splinter lemma into an infinite setting. In order to be able to obtain a statement like Lemma 10 in the infinite setting, we will need to 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 minimize the number of separations crossing them, an idea which also appeared in Carmesin's original proof of Theorem 4.4.1 in [9], as well as in [14] and our proof of the canonical splinter Lemma 11 for finite separation systems. However, our lemma 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 4.4.1 for arbitrary graphs, we will state our lemma in more generality here: not as a lemma about nestedness and separations, but as a lemma about a general nestedness-like relation, similar to the version Lemma 12 of our splinter lemma. This allows us to apply the lemma in Section 4.4.4 not to separations directly, where it would fail, but to substitute separators as a proxy giving our Theorem 16. From this result we will retrieve the separations for our proof of Theorem 4.4.1 in Section 4.4.4, but we will also build from this a tree of tree-decompositions to deduce Theorem 4.4.2 in Section 4.4.4.

The statement of our Lemma 15 is also inspired by our canonical splinter Lemma 11 for the finite setting, 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 2.3.1 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 4.4.4.

Now let $(\, \mathcal{A}_i : i \in I \,)$ be a family of non-empty subsets of $\mathcal{A}$ and $|\ |\colon I \to \mathbb{N}_0$ 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 properties:

(ST1) For every $i \in I$ all elements of $\mathcal{A}_i$ have finite k-crossing number for all $k \leqslant |i|$.

(ST2) 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.

(ST3) 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 15. *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 $N_k := N_{k-1} \cup N_k^+$ 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 (ST1) 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 (ST2), $\mathcal{A}_i$ 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 (ST3) 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^+. □

4.4.4 Applications of the thin splinter lemma

In this section we are going to apply Lemma 15 to infinite graphs. The application to locally finite graphs in Section 4.4.4 will be a straightforward application to a universe of separations, whereas in Section 4.4.4 we are going to use a more involved argument.

For either case we will utilize the fact that separations which efficiently distinguish two regular profiles are tight. Recall that 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 4.4.9. *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

$$(A \cap B, V) \vee (\bigcup \mathcal{C}_1 \cup Y_1, V \setminus \bigcup \mathcal{C}_1) \vee \cdots \vee (\bigcup \mathcal{C}_l \cup Y_l, V \setminus \bigcup \mathcal{C}_l) \in P.$$

However, for $l = m$ this contradicts the assumption since

$$(A \cap B, V) \vee (\bigcup \mathcal{C}_1 \cup Y_1, V \setminus \bigcup \mathcal{C}_1) \vee \cdots \vee (\bigcup \mathcal{C}_m \cup Y_m, V \setminus \bigcup \mathcal{C}_m)) = (B, A) \notin P\,.$$

□

Locally finite graphs

In this section we apply Lemma 15 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 15 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 (ST1) 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 $\mathcal{A}_{P,P'}$, (ST1) will follow immediately from the following assertion:

Proposition 4.4.10. *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 4.4.10 from the following lemma about tight separations:

Lemma 4.4.11. *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 $A' \smallsetminus B'$ and $B' \smallsetminus A'$ 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 $A \smallsetminus B$. Since $C' \subseteq A \smallsetminus B$ has no neighbours in $B \smallsetminus A$ 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 $B \smallsetminus A$. From $(A' \cap B') \subseteq A$ it follows that C does not meet $A' \cap B'$ and is hence contained in either $A' \smallsetminus B'$ or $B' \smallsetminus A'$. By possibly switching the roles of A' and B' we may assume that $C \subseteq A' \smallsetminus B'$. As above we can conclude from the tightness of C that $(A \cap B) \subseteq C$.

It remains to check two cases. If $(B \smallsetminus A) \cap (B' \smallsetminus A')$ 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 $(B \smallsetminus A) \cap (B' \smallsetminus A')$ 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 $N((B \smallsetminus A) \cap (B' \smallsetminus A')) \subseteq (A \cap B) \cap (A' \cap B')$. 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 $A' \smallsetminus B'$ and in $B' \smallsetminus A'$, 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 4.4.11 to establish Proposition 4.4.10:

Proof of Proposition 4.4.10. Since G is locally finite, the set $(A \cap B) \cup N(A \cap B)$ is finite. Therefore, by Lemma 4.4.4, 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 4.4.11 since we know by Lemma 4.4.9 that (A, B) is tight. □

The family $(\, \mathcal{A}_i : i \in I \,)$ therefore satisfies (ST1). With regard to (ST2) we can prove the following:

Lemma 4.4.12. *Let $(A, B) \in \mathcal{A}_{P,P'}$ and $(C, D) \in \mathcal{A}_{Q,Q'}$ with the property that $|(A, B)| < |(C, D)|$. Then some corner separation of (A, B) and (C, D) lies in $\mathcal{A}_{Q,Q'}$.*

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

$$|(A, B) \vee (C, D)| \leqslant |(C, D)| \text{ or } |(A, B) \vee (D, C)| \leqslant |(C, D)|,$$

it follows that this corner separation efficiently distinguishes Q and Q' by Lemma 2.3.1, so suppose that this is not the case. Then submodularity implies that

$$|(B, A) \vee (C, D)| < |(A, B)| \text{ and } |(B, A) \vee (D, C)| < |(A, B)|,$$

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'. $\square$

Corollary 4.4.13. *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. This is the assertion of Lemma 4.4.12. $\square$

It remains to show that $(\mathcal{A}_i : i \in I)$ satisfies (ST3). For this we need the following lemma:

Lemma 4.4.14. *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 $(C,D) \in \mathcal{A}_{P,P'}$ and Lemma 4.4.6 results in $(A,B) \vee (C,D) \in \mathcal{A}_{P,P'}$ and $(B,A) \vee (D,C) \in \mathcal{A}_{P,P'}$. Thus, by (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 4.4.6 and the disagreement of the four profiles on $\{A,B\}$ and $\{C,D\}$. $\square$

Using this lemma, we can now show that $(\mathcal{A}_i : i \in I)$ satisfies (ST3) using Lemma 4.4.5:

Lemma 4.4.15. *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 4.4.14 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 4.4.5 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. □

We are now ready to prove the main result of this subsection, which is similar to [14, Theorem 7.5]:

Theorem 4.4.16. *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 $\alpha : G \to G'$ is an isomorphism, then $\alpha(\mathcal{N}(G,\mathcal{P})) = \mathcal{N}(\alpha(G), \alpha(P))$.*

Proof. The combination of Proposition 4.4.10, Corollary 4.4.13, and Lemma 4.4.15 shows that the family $(\mathcal{A}_i : i \in I)$ thinly splinters. The nested set $N \subseteq \mathcal{A}$ produced by Lemma 15 meets each set $\mathcal{A}_i$ and thus distinguishes all pairs of distinguishable profiles in $\mathcal{P}$ efficiently. □

The nested set found by Theorem 4.4.16 does not in general correspond to a tree-decomposition of G, as Example 4.4.8 demonstrated. However, Theorem 4.4.16 can 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 4.4.7. In particular we can use Theorem 4.4.16 together with Lemma 4.4.7 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 distinguishable robust regular profiles, pairwise distinguishable by a separation of order at most k, a canonical tree-decomposition of G that efficiently distinguishes all profiles from $\mathcal{P}$.

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 4.4.10: 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 [33] gave an example of a graph which does not contain a canonical nested set of separations separating its ends. Since 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 different *separators* are involved, and then used that every separator appears in only finitely many separations. Thus, in this section instead of applying Lemma 15 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 4.4.1 (from [9]) 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 [9] and [14]. Moreover, the first step of this process also works for non-principal but regular profiles, allowing us to also get an (intermediate) result for those profiles, unlike the theorems from [9] and [14]. Note that distinguishing non-principal profiles is also discussed extensively in [45].

Many of the techniques applied throughout are similar to or inspired by arguments made in [14], particularly the approach of minimizing 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 4.4.17. *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 4.4.9, 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 (ST1) we need to make two basic observations about how the vertices of a crossing pair of separators in $\mathcal{A}$ lie:

Lemma 4.4.18. *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. □

Lemma 4.4.19. *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 4.4.18. Let v be a vertex in $Y \cap C_1$ and w a vertex in $Y \cap C_2$. Since both C_1 an C_2 are tight components, X is indeed a $\subseteq$-minimal v-w-separator. □

We can now combine these with Lemma 4.4.4 to show that we satisfy (ST1).

Lemma 4.4.20. *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 4.4.19, 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 4.4.4 every pair has only finitely many minimal separators of size k. Therefore, only finitely many such $Y \in \mathcal{A}$ exist. □

The following lemmas show how the separators of corner separations behave under our new nestedness relation. We will need these to show (ST2) and (ST3). Recall from Section 4.4.3 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 4.4.21. *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 4.4.9, a tight component C_X of $G - X$ which is disjoint from Z. By Lemma 4.4.18, 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 4.4.22. *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 4.4.9, 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 4.4.18, Z cannot cross the separator $(A_X \cup A_Y) \cap (B_X \cap B_Y)$. □

These now allow us to reuse the Lemmas 4.4.12 and 4.4.14 to show (ST2) and (ST3):

Lemma 4.4.23. *If two separators $X \in \mathcal{A}_{P,P'}$ and $Y \in \mathcal{A}_{Q,Q'}$ cross and $|P, P'| < |Q, Q'|$, 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 4.4.12 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 4.4.21 as desired. □

Lemma 4.4.24. *If two separators $X \in \mathcal{A}_{P,P'}$ and $Y \in \mathcal{A}_{Q,Q'}$ cross and $|P, P'| = |Q, Q'| = k$, 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 4.4.14 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 the Lemmas 4.4.21 and 4.4.22 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 15 are satisfied. Immediately we obtain the main result of this section:

Theorem 16. *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 4.4.1 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 [9]; 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 4.4.1, as Elm and Kurkofka [45, 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 16 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 4.4.25 ([45, 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 4.4.25 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:

Theorem 4.4.1 ([9, 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 4.4.25 to the set $\mathcal{N}$ of separators obtained from Theorem 16. 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 4.4.9 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. $\square$

For the reader's convenience, we also offer a direct proof of Theorem 4.4.1 which does not use Theorem 4.4.25. Instead, we perform an argument akin to one of the arguments used in the proof of Theorem 4.4.25 but in slightly simpler form, as the statement we need is a weaker one than Theorem 4.4.25.

Direct proof of Theorem 4.4.1. Let $\mathcal{N}$ be the nested set of separators obtained from Theorem 16 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:

Theorem 4.4.2 ([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:*

(a) the tree of tree-decompositions distinguishes $\mathcal{P}$ efficiently;

(b) if $t \in V(T)$ has level k, then $(T_t, \mathcal{V}_t)$ contains only separations of order k;

(c) 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 4.4.2 from Theorem 16 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 16 are strongly nested:

Lemma 4.4.26. *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 $Y \not\subseteq C \cup N(C)$, 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 4.4.9. 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. □

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 4.4.27. *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 $C_{Y'}$ so that $X' \subseteq C_{Y'} \cup N(C_{Y'})$ □

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 16. 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 4.4.28. *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 neither X nor Y. □

Now we are ready to deduce Theorem 4.4.2 from Theorem 16:

Proof of Theorem 4.4.2. 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-decomposition $(T_t, \mathcal{V}_t)$. We do this in a way that ensures the following properties:

(i) If $d(r,t) = k$, then every separation in $(T_t, \mathcal{V}_t)$ has order $k+1$.

(ii) 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 $d(t,r) \leqslant k$.

(iii) 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 4.4.28 gives that the separations are nested.

Moreover, every ω-chain $(A_1, B_1) < (A_2, B_2) < \dots$ 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 by construction, it is a regular tree set. Thus, by Lemma 4.4.7 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 (i) to (iii).

Observe that this decomposition satisfies (i) and (iii) 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, (ii) 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, for $k \geqslant 1$, 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 $G_{t'} \coloneqq G'$. 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 (ii) 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 4.4.28 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 (iii), 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 $Z \coloneqq N_{G_{t'}}(\bigcap_{i \in \mathbb{N}} B_i)$ 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 such 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, thus S_{k+1} is a regular tree set, so by Lemma 4.4.7 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 (i) to (iii). (i) is obvious. For (ii) 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 (iii) 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 component of $G_{t'} - X$. Hence, H meets at most one component of $G - X$. This gives (iii).

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 (a) to (c) from the assertion. The properties (b) and (c) are fulfilled by construction, so we only need to verify (a).

Let P, P' be two robust principal profiles from $\mathcal{P}$. By Theorem 16, $\mathcal{N}'$ 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. Since 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 $(T_t, \mathcal{V}_t)$. $\square$

4.5 Edge blocks

4.5.1 Introduction

In this section we are going to apply the results from Section 4.4 to another type of 'highly connected' substructures of an infinite graph: those which we think of as 'edge-connected'. Related results can be found in [8, 10–14, 23, 26, 27, 29–31, 33, 36, 39, 42, 45, 46, 48, 49, 53, 58, 68, 74].

Unlike for vertex-connectivity, where there are multiple competing notions of 'k-connected piece', i.e. tangles, blocks and profiles, for edge-connectivity, there does exist a single notion of 'k-edge-connected pieces' that undeniably is the most natural one. Let $k \in \mathbb{N} \cup \{\infty\}$ and let G be any connected graph, possibly infinite. We say that two vertices or ends are ($<k$)-*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. In this section of this thesis we 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. One main result of this section reads as follows:

Theorem 17. *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 [77], and they are more general than the 'tree-partitions' introduced by Seese [69] and by Halin [57]; see Section 4.5.4.

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 4.5.10) 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 17 true can be characterized in terms of generating bonds (for the definition of *generate* see Section 4.5.5):

Theorem 18. *Let G be any connected graph and let M be any nested set of bonds of G. Then the following assertions are equivalent:*

1. *M efficiently distinguishes all the edge-blocks of G;*
2. *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 18 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 18 is straightforward, Theorem 17 can be deduced from their theory, but it is not stated in [17] explicitly. Our Theorem 18 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 17 and the forward implication (i)→(ii) of Theorem 18 to obtain a purely combinatorial proof of the main result of Dicks and Dunwoody. In particular, this result shows again that Lemma 15 is highly flexible and can be used to distinguish highly connected structures in a variety of contexts. Together, our proofs of Theorem 17 and Theorem 18 take just over 7 pages in total.

This section is organized as follows. In Section 4.5.2 we introduce the tools and terminology that we need. In Section 4.5.3 we prove Theorem 17, and we show that we obtain a canonical set N. In Section 4.5.4 we relate each N_k to a tree-cut decomposition. In Section 4.5.5 we prove Theorem 18.

4.5.2 Tools and terminology

Throughout this section, $G = (V, E)$ denotes any connected undirected graph, finite or infinite. When we say ends we mean vertex-ends as usual, 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 following lemma is well-known [20, Exercise 8.12]; we provide a proof for the reader's convenience.

Lemma 4.5.1. *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 $B_0, B_1, \ldots$ of size k, say. Let F be an inclusion-wise 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. □

Corollary 4.5.2. *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 4.5.1. □

Cuts, bonds and separations

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 precisely if they are nested in the universe of bipartitions of V, i.e. if one of the four inclusions $A \subseteq C$, $A \subseteq D$, $B \subseteq C$ or $B \subseteq D$ holds.

Clearly, the set of all cut-separations forms a universe of separations (except for the fact that this universe needs to contain the empty cut $\{\emptyset, V\}$ which we not treat as a valid cut-separation), and consequently the set of all those cut-separations of order less than k, for some integer k, corresponds to a separation system S_k.

4.5.3 Proof of Theorem 17

The proof of our first main result will use Lemma 15, for that we need to show that we find the required corners for the corresponding nestedness relation $\sim$ on the set $\mathcal{A}$ of all the bond-separations of a connected graph G that efficiently distinguish some edge-blocks of G.

When $a = \{A, B\}$ and $b = \{C, D\}$ are two bond-separations, then we will consider as corners of a and b one of the following four possible objects: either $\{A \cap C, B \cup D\}$, $\{A \cap D, B \cup C\}$, $\{B \cap D, A \cup C\}$ or $\{B \cap C, A \cup D\}$. These are the four possibilities of how a new cut-separation can be built from $\{A, B\}$ and $\{C, D\}$ using just '$\cup$' and '$\cap$' and also correspond to the corner separations of $\{A, B\}$ and $\{C, D\}$ in the universe of cut-separations. Note that sometimes an intersection may be empty so some of the four possibilities may not be valid cut-separations; and sometimes a possibility is a cut-separation but not an element of $\mathcal{A}$. We will see in Lemma 4.5.4 that every possibility that happens to lie in $\mathcal{A}$ is already a corner of $\{A, B\}$ and $\{C, D\}$ in the sense of our splinters hierarchically condition, provided that $\{A, B\}$ and $\{C, D\}$ cross.

As the index set I of our family $(\mathcal{A}_i : i \in I)$ of non-empty subsets of $\mathcal{A}$ to which we want to apply Lemma 15 we will consider, as usual, 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. Then every $\mathcal{A}_i$ will be non-empty because the edge-blocks forming i are disjoint. Our choice for the order $|i|$ of an index i shall 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 i will be a k-edge-block for some $k > |i|$.

Now in order to employ Lemma 15 to deduce Theorem 17, we first have to verify that $(\, \mathcal{A}_i \,:\, i \in I\,)$ thinly splinters. To this end, we verify (ST1) to (ST3) below.

Recall that the *k-crossing number* of a, for an $a \in \mathcal{A}$ and $k \in \mathbb{N}$, is the number of elements of $\mathcal{A}$ that cross a and lie in some $\mathcal{A}_i$ with $|i| = k$. Note that in our case, every bond-separation of order k can only possibly lie in sets $\mathcal{A}_i$ with $|i| = k$. Thus, the k-crossing number of a bond-separation of arbitrary finite order will be the number of efficiently distinguishing bond-separations of order k crossing it. Thus, we can deduce that every k-crossing number is finite, and thus that (ST1) holds, from Lemma 4.5.1:

Lemma 4.5.3. *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 pigeonhole 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 pigeonhole principle, infinitely many of these bonds must contain the same edge of P. This contradicts Lemma 4.5.1. □

Next, to show (ST2), we need the following lemma:

Lemma 4.5.4. *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).*

Proof. Since $\{X, Y\}$ is a cut-separation that is nested with both, $\{A_1, B_1\}$ and $\{A_2, B_2\}$, either X or Y is a subset of B_1 or B_2, in which case it is immediate that $\{X, Y\}$ is nested with $\{A_1 \cap A_2, B_1 \cup B_2\}$ as desired, or, one of X and Y is a subset of A_1 and one of X and Y is a subset of A_2. However, since $A_1 \cup A_2 \neq V(G)$ (as $\{A_1, B_1\}$ and $\{A_2, B_2\}$ cross) it needs to be the case that either $X \subseteq A_1 \cap A_2$ or $Y \subseteq A_1 \cap A_2$, so in either case $\{X, Y\}$ is nested with $\{A_1 \cap A_2, B_1 \cup B_2\}$ as desired. □

Using this lemma, we can now show (ST2):

Lemma 4.5.5. *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 4.5.4, 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. However, this follows as the order function on the set of bond-separations is submodular: 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 (ST3), we need the following lemma:

Lemma 4.5.6. *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 (ST3):

Lemma 4.5.7. *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, by submodularity,

$$|E(A \cap C, B \cup D)| + |E(B \cap D, A \cup C)| \leqslant |E(A, B)| + |E(C, D)| = 2k, \quad (1)$$

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 4.5.4. 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 (ST1), and the two cut-separations $\{A, B\}$ and $\{C, D\}$ cross which allows us to deduce the desired inequality between the sums by the Lemmas 4.5.4 and 4.5.6, as follows:

- by Lemma 4.5.4, 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 4.5.6, every $\{X, Y\} \in \mathcal{A}$ of order k that crosses both cut-separations $\{A \cap C, B \cup D\}$ 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 (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 4.5.4. 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. □

We can now prove the first main result of this section:

Proof of Theorem 17. Let G be any connected graph. By the Lemmas 4.5.3, 4.5.5, and 4.5.7 we may apply Lemma 15 to the family $(\, \mathcal{A}_i \,:\, i \in I \,)$ defined at the beginning of this 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 15, we indeed obtain that $\varphi(N(G)) = N(\varphi(G))$. □

4.5.4 Nested sets of bonds and tree-cut decompositions

Recall that, given a connected graph G, we denote by $N = N(G)$ the canonical set of nested bonds from Theorem 17 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 that:

- 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. In this section of this thesis, 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 4.5.8. 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 $v_i v_j$ $(i \neq j \in [3])$. 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 such 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.

Similar like a tree-decomposition forms an S-tree, so does a tree-cut decomposition: a tree-cut decomposition $(T, \mathcal{X})$ makes T into an S-tree for the set S of cut-separations which correspond to its fundamental cuts. We will use this observation to relate the set N_k to a tree-cut decomposition, by using Theorem 2.6.1 by Gollin and Kneip.

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 4.5.9. *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 2.6.1 to find an S-tree (T, α) so that the map $\alpha\colon \vec{E}(T) \to \vec{S}$ is an isomorphism. Since clearly no separation in $\vec{S}$ is small and so S is a regular tree set, for this, it suffices to show that B cannot contain pairwise distinct bonds $F_0, F_1, \ldots, F_\omega$ 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 4.5.2.

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 4.5.2 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 4.5.10. This example is a variation of Example 4.4.8. Consider the locally finite graph displayed in Fig. 4.5. 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 $w_1^n, \ldots, w_{n+2}^n$. Then we select 2^n vertices of the $K^{2^{n+2}}$ and call them $u_1^n, \ldots, u_{2^n}^n$. Furthermore, we select 2^{n+1} vertices of the $K^{2^{n+2}}$, other than the previously chosen u_i^n, and call them $v_1^n, \ldots, v_{2^{n+1}}^n$. 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.

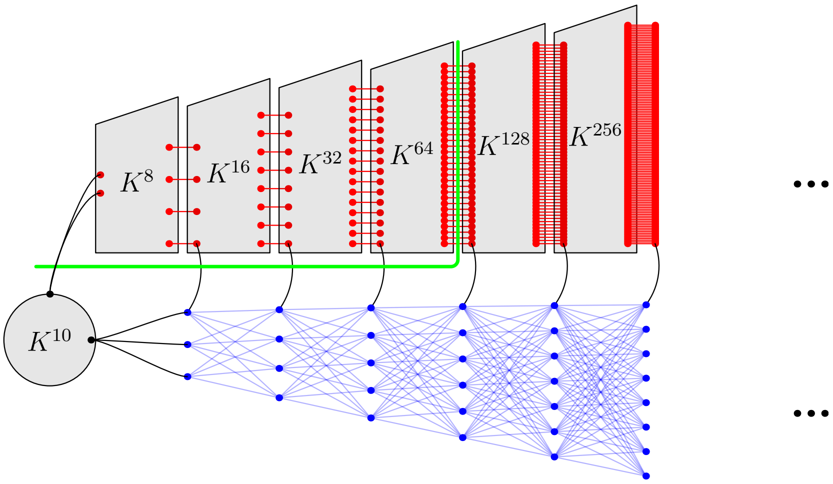

Figure 4.5: The only cut that efficiently distinguishes the two edge-blocks defined by K^{64} and by K^{128} is drawn in green.

Now the vertex sets of the chosen $K^{2^{n+2}}$ are $(2^{n+2}-1)$-edge-blocks B_n. The only cut-separation that efficiently distinguishes B_n and B_{n+1} is

$$F_n := \{ \bigcup_{k=1}^{n} B_n \, , \, V \setminus \bigcup_{k=1}^{n} B_n \, \}.$$

Additionally, the vertex set of the K^{10} is a 9-edge-block B_0. The only cut-separation that efficiently distinguishes 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 2.6.1.

4.5.5 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, where

- $(A, B) \vee (A', B') := (A \cup A', B \cap B')$ is the *supremum* and
- $(A, B) \wedge (A', B') := (A \cap A', B \cup B')$ is the *infimum*

of two cut-separations (A, B) and (A', B'). In this section we prove Theorem 18:

Theorem 18. *Let G be any connected graph and let M be any nested set of bonds of G. Then the following assertions are equivalent:*

1. *M efficiently distinguishes all the edge-blocks of G;*

2. *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 need a generalized version of the star-comb lemma [20, Lemma 8.2.2]. A *comb* in a given graph G means 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 4.5.11 (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 [20, 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 inclusion-wise minimal tree $T \subseteq G$ that extends all the rays R_ω with $\omega \in U$. Let $T' \subseteq T$ be the inclusion-wise 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. □

For more on stars and combs, see the series [4–7].

We are now ready to conclude this section of this thesis with a proof of Theorem 18:

Proof of Theorem 18. (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, as both X and Y are obtained as by a combination of unions and intersections of the A_ℓ and B_ℓ. This contradicts 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 4.5.9, 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 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 $(T, \mathcal{X})$. $\quad (*)$

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. We 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. We colour all the nodes in a component of this forest red if the component has a red neighbour, and blue otherwise.

We let $T_1 \subseteq T$ be the forest induced by the red nodes, and we 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 $s_0t_0, \dots, s_nt_n$ are the finitely many T_1–T_2 edges with $s_\ell \in T_1$ and $t_\ell \in T_2$, then

$$(V_1, V_2) = \bigwedge_{C\colon \text{ a component of } T_2} \quad \bigvee_{\ell\colon\, t_\ell \in C} \alpha(s_\ell, t_\ell) \;.$$

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 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 all the sets $Y_c := \bigcup \{\, \hat{X}_t : t \in c \,\}$ 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, a contradiction. $\square$

4.6 Obtaining trees of tangles from tangle-tree duality

4.6.1 Introduction

In this section of this thesis we demonstrate the versatility of the most abstract version of the tangle-tree duality theorem: we deduce Theorem 1.1.3 and some of its variations from it, 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 1.1.3, we exploit the generality of the most abstract version of the tangle-tree duality theorem, which reads as follows:

Theorem 4.6.1 (Tangle-tree duality theorem [30, 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 $\mathcal{F}$.*

The strength of Theorem 4.6.1 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 4.6.1 to be employed in a multitude of different settings ([26, 29]). The freedom in choosing and manipulating $\mathcal{F}$ will also allow us to achieve our goal of deducing tree-of-tangles theorems from Theorem 4.6.1: 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 present multiple variations of this idea throughout this section of this thesis.

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 in [27],[26] or Section 4.1, our general purpose solution to obtaining tree-of-tangles theorems in a wide range of structures. (There, we showed an even more general theorem than Theorem 1.1.3 which no longer mentions tangles or profiles at all, but just talks about sets of separations fulfilling one simple-to-check condition.)

Instead of competing in terms of simplicity and brevity just for a proof of the tree-of-tangles theorem, the aim of this section 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 [26, 27, 39].

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 4.6.1 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 4.6.1 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.

The structure of this section is as follows. In Section 4.6.2 we will repeat the required definitions from [21, 26, 27, 29, 30] not already given in Chapter 2. In Section 4.6.3 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 4.6.4, 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 4.6.5 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 4.6.6. In Section 4.6.7 we prove a tree-of-tangles theorem for tangles of different orders.

4.6.2 Terminology and background

The tree-of-tangles theorem (see [27])

For the purpose of this section, we shall use the following slightly more restrictive, non-canonical version of Theorem 1.1.2:

Theorem 4.6.2 ([27, Corollary 3.7], modified)**.** *Let $(\vec{U}, \leqslant, ^*, \vee, \wedge, |\ |)$ be a submodular universe of separations. For every set $\mathcal{P}$ of pairwise distinguishable robust regular profiles in $\vec{U}$ there is a regular tree set $T = T(\mathcal{P}) \subseteq \vec{U}$ of separations such that:*

(i) every two profiles in $\mathcal{P}$ are efficiently distinguished by some separation in T;

(ii) every separation in T efficiently distinguishes a pair of profiles in $\mathcal{P}$.

Apart from the nested set T not being canonical, the major difference between Theorem 1.1.2 and Theorem 4.6.2 lies in the fact that we require here that the set $\mathcal{P}$ of profiles is regular and robust, whereas Theorem 1.1.2 only required $\mathcal{P}$ to be a robust set of profiles. Moreover, we do not directly achieve that T is a regular tree set, this however can easily be fixed at a later point: if T is a nested set of separations which efficiently distinguish any two regular profiles from $\mathcal{P}$, then just deleting all separations with a small orientation from T results in a regular tree set doing so as well, as no small separation can distinguish two regular profiles.

Tangle-tree duality (see [30])

Given some set $\mathcal{F}$ of subsets of S, an $\mathcal{F}$*-tangle* of S is a consistent orientation of S which includes no subset in $\mathcal{F}$. Given a submodular universe $\vec{U}$, we say that τ is an $\mathcal{F}$*-tangle in* $\vec{U}$ if τ is an $\mathcal{F}$-tangle of some S_k. Observe that profiles are $\mathcal{F}_P$-tangles for the set $\mathcal{F}_P$ of all 'profile triples' $\{\vec{r}, \vec{s}, (\vec{r} \vee \vec{s})^*\} \subseteq \vec{S}$.

Often we will consider sets $\mathcal{F}$ of stars.

We say that a set $\mathcal{F}$ *forces* a separation $\vec{s} \in \vec{S}$ if $\{\overleftarrow{s}\} \in \mathcal{F}$.

$\mathcal{F}$ is *standard for* $\vec{S}$ if it forces all trivial separations, that is $\mathcal{F}$ contains all singletons $\{\overleftarrow{s}\}$ for co-trivial $\vec{s} \in \vec{S}$.

Recall that an S-tree (T, α) consists of a tree T together with a map α from $\vec{E}(T)$ to $\vec{S}$ which commutes with *. Given some set $\mathcal{F}$ of subsets of S, such an S-tree (T, α) is *over* $\mathcal{F}$ if $\alpha(t) \in \mathcal{F}$ for all $t \in V(T)$, i.e. for very node t of T the set of all ingoing separations at that node is contained in $\mathcal{F}$.

An S-tree (T, α) is *irredundant* if for any 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 any 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, for every $\vec{s} \geqslant \vec{r}$,

$$f\!\downarrow^{\vec{r}}_{\vec{s}_0}(\vec{s}) = \vec{s} \vee \vec{s}_0 \quad \text{and} \quad f\!\downarrow^{\vec{r}}_{\vec{s}_0}(\overleftarrow{s}) = (\vec{s} \vee \vec{s}_0)^*.$$

This map is defined on $\vec{S}_{\geqslant \vec{r}} \smallsetminus \{\overleftarrow{r}\}$, where $S_{\geqslant \vec{r}}$ is the set of all separations $t \in S$ which have an orientation $\vec{t}$ with $\vec{t} \geqslant \vec{r}$, and $\vec{S}_{\geqslant \vec{r}}$ is the set of all orientations of separations in $S_{\geqslant \vec{r}}$.

For an irredundant S-tree (T, α) over some set of stars with $\{\overleftarrow{r}\} = \alpha(x)$, for some leaf x of T, we write

$$\alpha_{x, \vec{s}_0} \coloneqq f\!\downarrow^{\vec{r}}_{\vec{s}_0} \circ \alpha\,.$$

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

$$\sigma^{\vec{s}_0}_{\vec{x}} \coloneqq \{\vec{x} \vee \vec{s}_0\} \cup \{\vec{y} \wedge \overleftarrow{s}_0 \ : \ \vec{y} \in \sigma \smallsetminus \{\vec{x}\}\}\,,$$

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 $\vec{t} \in \vec{S} \smallsetminus \{\overleftarrow{r}\}$ 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 separation system S is *separable* if for any two non-trivial nondegenerate 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 nondegenerate 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 4.6.1 (Tangle-tree duality theorem [30, 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 4.6.1.

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 [26, Lemma 12].)

We shall need the following additional lemmas from the literature:

Lemma 4.6.3 ([30, 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 4.6.4 ([30, 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 4.6.5 ([26, 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 [30, Lemma 4.2] which follows with the exact same proof:

Lemma 4.6.6 ([30])**.** *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 nondegenerate, 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 4.6.6 and [30, Lemma 4.2] is that [30, 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 [30, 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 $\alpha(t) \in \mathcal{F}$.

Splices in submodular universes

In addition to the existing terminology, we shall need the following new concept, which has already been considered in [29], but has not been 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 [29] which, phrased in our terminology, considers a splice between two separations. We recapitulate the main argument of this proof below.

Lemma 4.6.7 ([29])**.** *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 [29, 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 2.3.1, $\vec{r} \leqslant \vec{t} \wedge \vec{s} \leqslant \vec{s}$ and this contradicts the fact that $\vec{s}$ is a splice for $\vec{r}$. □

This lemma then directly implies the ultimate statement of [29, Lemma 3.4]:

Lemma 4.6.8 ([29, Lemma 3.4])**.** *Every $\vec{S}_k \subseteq \vec{U}$ in a submodular universe is separable.*

4.6.3 Structurally submodular separation systems

In this section we will prove the first tree-of-tangles theorem of this section of this thesis. It is a theorem for regular profiles, all of the same structurally submodular separation system, and states as follows:

Theorem 4.6.9. *Let $\vec{S}$ be a structurally submodular separation system. Then S contains a nested set that distinguishes the set of regular profiles of S.*

By itself Theorem 4.6.9 is nothing special; indeed, it is a slight weakening of Theorem 1.1.3, 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 4.6.9 as a direct consequence of Theorem 4.6.1.

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 4.6.1 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 stars and if we could feed this $\mathcal{F}$ to Theorem 4.6.1, 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 $\mathcal{F}$ is 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 $\max P \in \mathcal{M}$ is included only in P itself.

Unfortunately, we are still some way off from plugging $\mathcal{F}$ into Theorem 4.6.1: 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 4.6.5.

We thus need to show that S is not only separable but $\mathcal{F}$-separable. Unfortunately our current set $\mathcal{F}$ is not even a set of stars yet. However, in [23] a solution was laid out for this exact situation: a series of lemmas from [23] shows that we can simply *make* $\mathcal{F}$ 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 4.6.10 ([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 4.6.11 ([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 4.6.1. However, the proof of Lemma 4.6.11 does not necessitate this, and neither does its application.

Lemma 4.6.11 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 $\mathcal{G}_{i-1}$. Write $\hat{\mathcal{F}}^* := \bigcup_{i \in \mathbb{N}} \mathcal{G}_i$. Then, by Lemma 4.6.11, the $\hat{\mathcal{F}}^*$-tangles of S are precisely its $\mathcal{F}^*$-tangles, which is to say that there are no $\hat{\mathcal{F}}^*$-tangles of S. Moreover, this set $\hat{\mathcal{F}}^*$ still has the property that each star in it is included in at most one regular profile: let us say that $\hat{\sigma^*} \in \hat{\mathcal{F}}^*$ *originates from* $\sigma \in \mathcal{F}$ if $\hat{\sigma^*}$ can be obtained by a series of shifts from an uncrossing of σ. The Lemmas 4.6.10 and 4.6.11 then say that, if $\hat{\sigma^*} \subseteq P$ for a regular profile P, and $\hat{\sigma^*}$ 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 $\hat{\sigma^*}$.

We can thus formally prove Theorem 4.6.9:

Proof of Theorem 4.6.9. 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 4.6.5 S is separable. Together this gives that S is $\mathcal{F}$-separable. Hence, we can apply the tangle-tree duality Theorem 4.6.1 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 the Lemmas 4.6.10 and 4.6.11 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 4.6.10 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 4.6.1, which we may assume to be irredundant (Lemma 4.6.4). Let $\vec{N}$ be the image of α. Then N is a nested subset of S (Lemma 4.6.3). 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. □

Let us make some remarks on this proof of Theorem 4.6.9. 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 4.6.4, 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 4.6.9 to the level of Theorem 1.1.3 by dropping the assumption of regularity, since Lemma 4.6.11 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 4.6.5 and 4.6.7.

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 4.6.1 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 4.6.11; 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 4.6.10.

Lemma 4.6.12. *Let $\vec{S} \subseteq \vec{U}$ be a structurally submodular separation system 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 2.3.1, 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

$$\sigma' := \sigma \smallsetminus \{\vec{t}\} \cup \{\vec{t} \wedge \vec{s},\, \vec{t} \wedge \overleftarrow{s}\}$$

also includes σ. Consequently, σ' together with s are a better choice than σ and s, a contradiction. $\square$

Lemma 4.6.13. *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 $\alpha : N' \to N$ 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 $N \smallsetminus \{\vec{s}\} \cup \{\vec{s} \wedge \overleftarrow{t}\}$ is also a candidate for N, contradicting the $\leqslant$-minimality. So suppose that $N \smallsetminus \{\vec{s}\} \cup \{\vec{s} \wedge \overleftarrow{t}\}$ 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 4.6.9 without resorting to Lemma 4.6.10:

Theorem 4.6.9. *Let $\vec{S}$ be a structurally submodular separation system. Then S contains a nested set that distinguishes the set of regular profiles of S.*

Direct Proof. 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 4.6.11, 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 co-small separations are contained in no regular profile.

By Theorem 4.6.1 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 4.6.13 every consistent orientation which avoids $\mathcal{F}_{\mathcal{P}}$ is a profile and if P would not be regular, it would contain a co-small separation $\vec{s}$ which is impossible, since $\{\vec{s}\} \in \mathcal{F}_{\mathcal{P}}$. So, by Lemma 4.6.12, 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. $\square$

4.6.4 Application: Degrees in trees of tangles

In this section we are going to see that our proof of Theorem 4.6.9 in Section 4.6.3 has one advantage over the usual, more direct proofs of Theorem 4.6.9 from [26] and Section 4.1: it allows us to easily control the maximum degree of the resulting tree. More precisely: let S be a structurally submodular separation system 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 4.6.14. *Let S be a structurally submodular separation system. 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 4.6.14 we will follow the first proof of Theorem 4.6.9, 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 structurally 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 4.6.14 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 4.6.14 is at most $|\max P|$.

Let us now prove Theorem 4.6.14:

Proof of Theorem 4.6.14. 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 4.6.9, 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 4.6.9 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 4.6.9 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 $\Delta(T) = \delta(P)$. □

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 $\delta(P) \leqslant 2$:

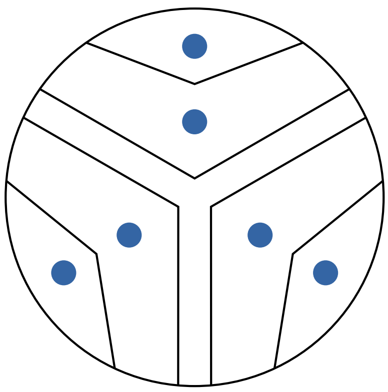

Figure 4.6: A ground-set and system of bipartitions.

Example 4.6.15. Let V consist of the six points in Fig. 4.6, and S be the separation system given by the six outlined bipartitions of V together with $\{\emptyset, V\}$. (That is, $\vec{S}$ contains (A, B) and (B, A) for each of these bipartitions $\{A, B\}$. We have $(A, B) \leqslant (C, D) :\Leftrightarrow A \subseteq C$, and $(A, B)^* = (B, A)$. Compare [29].) 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.

4.6.5 Efficient distinguishers

Often our structurally submodular separation system $\vec{S}$ is actually an $\vec{S}_k$, the set of all separations of order less than k, of some submodular universe $\vec{U}$.

In this case we are not just interested in a nested set N of separations which distinguishes all profiles, but one which does so *efficiently*. Recall that this means that N, for any two profiles, 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 $\vec{S}_k$ utilizing the duality theorem together with a separate application of its core mechanism: shifting S-trees.

We will prove this theorem:

Theorem 4.6.16. *Let $\vec{U}$ be a submodular universe and let $\mathcal{P}$ be a set of regular profiles of $\vec{S}_k$. 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 4.6.3, 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* $\mathrm{Eff}(P)$ if the following holds:

$$\nexists\, \vec{s} \in \sigma \text{ and } \vec{s'} \in P\colon \vec{s} \leqslant \vec{s'} \text{ and } |s'| < |s|. \qquad (\mathrm{Eff}(P))$$

This condition ensures that, for two profiles P and P', a star σ with property $\mathrm{Eff}(P)$ containing $\vec{s}$, and a star σ' with property $\mathrm{Eff}(P')$ 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 $\mathrm{Eff}(P)$ whereas $\vec{r} \leqslant \vec{s}$ would contradict property $\mathrm{Eff}(P')$. 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 $\mathrm{Eff}(P)$ or $\mathrm{Eff}(P')$, respectively.

Property $\mathrm{Eff}(P)$ is preserved under taking shifts:

Lemma 4.6.17. *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 $\sigma' := \sigma_{\vec{x}}^{\vec{s}}$ and σ has property $\mathrm{Eff}(P)$, then also σ' has property $\mathrm{Eff}(P)$.*

Proof. Suppose for a contradiction that σ' does not have property $\mathrm{Eff}(P)$, 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 $|\vec{s} \wedge v| \geqslant |\vec{s}|$, 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 $\mathrm{Eff}(P)$. If however $\vec{t'}$ crosses $\vec{t}$, then, by the profile property of P and property $\mathrm{Eff}(P)$ 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 $|v' \wedge \vec{s}| \geqslant |\vec{s}|$, 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 $\mathrm{Eff}(P)$.

From the Lemmas 4.6.11 and 4.6.17 immediately we obtain the following corollary:

Corollary 4.6.18. *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 4.6.16 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 $\mathrm{Eff}(\cdot)$, 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 4.6.16 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 4.6.12, we define the *fatness* of a star σ as the tuple $(n_{k-1}, n_{k-2}, \dots, n_1, n_0)$, 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 4.6.19. *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 4.6.12 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 $\mathrm{Eff}(P)$. 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 4.6.16:

Proof of Theorem 4.6.16. We may apply Theorem 4.6.1 for $\mathcal{F}_e$ since $\vec{S}_k$ is $\mathcal{F}_e$-separable by Corollary 4.6.18 and $\mathcal{F}_e$ is standard since co-trivial 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 4.6.19, and Lemma 4.6.13 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 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 2.3.1, there indeed needs to exist 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 Eff(P) or 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 $(T_{P_1}, \alpha \restriction T_{P_1})$ 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 4.6.6 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 of 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 4.6.7. 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 4.6.7. 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). □

4.6.6 Degrees in efficient trees of tangles

In this section we apply our method from Section 4.6.4 to Theorem 4.6.16 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 4.6.4, 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 $\mathrm{Eff}(P)$ 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 4.6.19, there exists such a star for every regular profile P, thus $\delta_e(P)$ is a well-defined natural number.

We denote by $\delta_{e,\max}$ 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 4.6.20. *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 $D_P \subseteq P$ so 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 $\mathrm{Eff}(P)$ 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. □

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 $\delta_{e,\max}$ we also need to limit the size of the stars in $\mathcal{F}_e$ which are contained in no profile to $\delta_{e,\max}$. Like in Section 4.6.4, we cannot limit the maximum degrees below 3. Along the lines of the proof of Lemma 4.6.10, 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 4.6.21. *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. Since 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 4.6.16, which shows that we can find a tree of tangles of bounded degree:

Theorem 19. *Let $\vec{U}$ be a submodular universe and let $\mathcal{P}$ be the set of regular profiles of $\vec{S}_k$. Then there exists a 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* $\max\{\delta_e(\mathcal{P}), 3\}$.

Proof. Let $\mathcal{F}_e^s$ 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 $\max\{\delta_e(\mathcal{P}), 3\}$ 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 4.6.18. Moreover, the shift of a star cannot contain any profile which does not contain the original star by Lemma 4.6.11, thus S_k is also $\mathcal{F}_e^s$-separable.

Thus, all we need to show is that applying Theorem 4.6.1 cannot result in an $\mathcal{F}_e^s$-tangle, the rest of the proof can then be carried out as the proof of Theorem 4.6.16: instead of S-trees over $\mathcal{F}_e$ we now consider S-trees over $\mathcal{F}_e^s$, and observe that the shifting argument in the proof of Theorem 4.6.16 again shifts stars in $\mathcal{F}_e^s$ to stars in $\mathcal{F}_e^s$.

However, applying Theorem 4.6.1 indeed cannot result in an $\mathcal{F}_e^s$-tangle: such a tangle cannot be a regular profile, since by our definition of $\delta_e(P)$, there is a star in $\mathcal{F}_e^s$ contained in P. But every consistent orientation which is not a regular profile either contains a star $\{\vec{s}\}$ for a co-small separation $\vec{s}$ – each such star is also contained in $\mathcal{F}_e$ – or contains, by Lemma 4.6.21 a star of size 3 not contained in any profile. Either such star is also contained in $\mathcal{F}_e^s$ by definition. □

4.6.7 Tangles of mixed orders

In this section we would like to use the ideas from Section 4.6.5 to obtain a proof of Theorem 4.6.2 using tangle-tree duality. The challenge of Theorem 4.6.2 compared to Theorem 4.6.16 is that the set of profiles $\mathcal{P}$ considered in Theorem 4.6.2 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 4.6.2, we need a slightly stronger version of Theorem 4.6.1, which allows us to exclude this specific situation in the requirement of $\mathcal{F}$-separability.

More precisely, we need the following result from [41]:

Theorem 4.6.22 ([41, Theorem 7.1.]). *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.*

Here, a separation $\vec{r}$ in $\vec{S}$ is said to be $\mathcal{F}$*-critical* if there exists a star $\sigma \in \mathcal{F}$ such that $\vec{r} \in \sigma$, but there is no $\sigma' \in \mathcal{F}$ with $\sigma' \cap r = \{\overleftarrow{r}\}$. S is said to be *critically $\mathcal{F}$-separable* if all $\mathcal{F}$-critical separations $\vec{r}, \overleftarrow{r'} \in \vec{S}$ with $\vec{r} \leqslant \overrightarrow{r'}$ satisfy that there exists a separation $\vec{s}_0 \in \vec{S}$ which emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$ such that $\overleftarrow{s}_0$ emulates $\overleftarrow{r'}$ in $\vec{S}$ for $\mathcal{F}$. In particular, every $\mathcal{F}$-separable separation system S is also critically $\mathcal{F}$-separable.

We can use this version of the duality statement to obtain a result similar to Theorem 4.6.2, however our construction only works in distributive universes – recall that this means that $\vec{r} \vee (\vec{s} \wedge \vec{t}) = (\vec{r} \vee \vec{s}) \wedge (\vec{r} \vee \vec{s})$, always – since we need the following result from [34], which can be also found in [27]:

Lemma 4.6.23 ([27, Theorem 3.11], [34, 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 exist 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 $\mathrm{Eff}(P)$.

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 4.6.24. *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 implies 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 4.6.23 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.

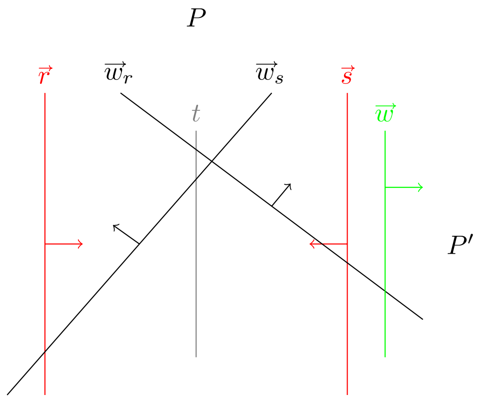

Similarly, if $\overleftarrow{w} \wedge \overrightarrow{w_r}$ had order less than the order of $\overrightarrow{w_r}$, this would contradict the choice of $\overrightarrow{w_r}$: by consistency P would need to contain $\overleftarrow{w} \wedge \overrightarrow{w_r}$ which contradicts the definition of M_r^P, from which $\overrightarrow{w_r}$ was chosen.

Thus, by submodularity, $\vec{w} \wedge \overrightarrow{w_r}$ 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 $\overleftarrow{w} \vee \overrightarrow{w_r}$ and therefore either weakly orients $\vec{s} \wedge \overrightarrow{w_r}$ as $\overleftarrow{s} \vee \overleftarrow{w_r}$ or $\vec{s} \wedge \overleftarrow{w_r}$ as $\overleftarrow{s} \vee \overrightarrow{w_r}$.

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 $\overrightarrow{w_r}$ does not cross $\vec{s}$ and, by a symmetric argument, that $\overrightarrow{w_s}$ does not cross $\vec{r}$. Hence, $\vec{r} \leqslant \overrightarrow{w_s}$ and $\vec{s} \leqslant \overrightarrow{w_r}$. We may therefore assume without loss of generality that $\overrightarrow{w_r} = \overrightarrow{w_s}$.

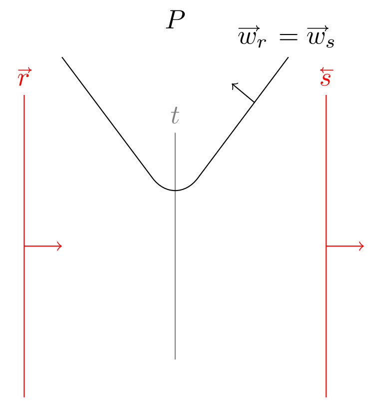

If $\overrightarrow{w_r} = \overrightarrow{w_s}$ crosses t, then, by the choice of t, neither $\overrightarrow{w_r} \wedge \vec{t}$ nor $\overrightarrow{w_r} \wedge \overleftarrow{t}$ has order less than $|t|$, thus $\overrightarrow{w_r} \vee \vec{t}$ and $\overrightarrow{w_r} \vee \overleftarrow{t}$ both have order at most the order of $\overrightarrow{w_r}$. By the strong robustness of P applied to $\overrightarrow{w_r}$, $\overleftarrow{w_r} \wedge \overleftarrow{t}$ and $\overleftarrow{w_r} \wedge \vec{t}$, we know that either $\overrightarrow{w_r} \vee \vec{t} \in P$ or $\overrightarrow{w_r} \vee \overleftarrow{t} \in P$. However, both contradict the $\leqslant$-maximal choice of $\overrightarrow{w_r}$. So, instead $\overrightarrow{w_r}$ is nested with t, that is, t has an orientation $\vec{t}$ such that $\vec{t} \leqslant \overrightarrow{w_r}$, so t is weakly oriented by P, as claimed. $\square$

Note that the assumption that our profiles are *strongly* robust is essential in this argument, for example for the case $\overrightarrow{w_r} = \overrightarrow{w_s}$: if we only assume robustness, we can not conclude that P contains either $\overrightarrow{w_r} \vee \vec{t}$ or $\overrightarrow{w_r} \vee \overleftarrow{t}$ and thus would not obtain a contradiction.

The next step is to verify that shifting with a splice as in Lemma 4.6.24 maps stars in $\mathcal{F}_d$ to stars in $\mathcal{F}_d$, which will prove that U is critically $\mathcal{F}_d$-separable:

Lemma 4.6.25. *Let r and s_0 be separations which are weakly oriented by every profile in $\mathcal{P}$ and suppose that $\vec{s}_0$ is 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 4.6.7, $\vec{s} \vee \vec{s}_0$ 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 4.6.11, 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 $\sigma^{\vec{s}_0}_{\vec{x}}$, and thus either P contains the inverse of some separation in $\sigma^{\vec{s}_0}_{\vec{x}}$ or $\sigma^{\vec{s}_0}_{\vec{x}} \subseteq P$.

Hence, by Lemma 4.6.17 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 $\overrightarrow{y'} \in \sigma^{\vec{s}_0}_{\vec{x}}$.

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 $\overleftarrow{w_y}$ also witnesses that P weakly contains $\overleftarrow{y} \wedge \overleftarrow{s_0}$ while $\vec{y} \vee \vec{s}_0 \in \sigma^{\vec{s}_0}_{\vec{x}}$. 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 $\overleftarrow{y} \vee \vec{s}_0 \in \sigma^{\vec{s}_0}_{\vec{x}}$.

Thus, we may suppose that P weakly orients s_0 as $\vec{s}_0$, witnessed by $\vec{w}_0 \in P$.

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, $\overleftarrow{s_0} \wedge \vec{w}_y$ has order at most the order of w_y, i.e. it is oriented by P. By Lemma 4.6.23 applied to $\vec{w}_0, \overleftarrow{w_y} \in P$ and $\overleftarrow{s_0} \wedge \vec{w}_y$ we can therefore conclude that P contains $\vec{s}_0 \vee \overleftarrow{w_y}$, i.e. P weakly contains $\overleftarrow{y} \vee \vec{s}_0 \leqslant \vec{s}_0 \vee \overleftarrow{w_y}$. □

In order to use our stronger tangle-tree duality Theorem 4.6.22 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 this in the following two lemmas.

Lemma 4.6.26. *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 $\vec{s}_Q$. □

Unlike for structurally submodular separation systems in Lemma 4.6.13 or efficient distinguishers in Lemma 4.6.21, 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 4.6.21, 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 4.6.22 does not result in an $\mathcal{F}_d$-tangle.

Lemma 4.6.27. *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 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 an 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 choice 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.

Hence, 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.

Therefore, σ contains precisely one profile and therefore, by construction, $\sigma \in \mathcal{F}_d$. $\square$

Together with Theorem 4.6.22, 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 4.6.28 (Tree-of-tangles theorem for different orders)**.** *Let $\vec{U} = (\vec{U}, \leqslant, ^*, \vee, \wedge, |\cdot|)$ 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 4.6.24 and Lemma 4.6.25 the set U is critically $\mathcal{F}_d$-separable for the set $\mathcal{F}_d$ defined above. Thus, we can apply Theorem 4.6.22. This can, by Lemma 4.6.27, not result in an $\mathcal{F}_d$-tangle, thus there is a U-tree over $\mathcal{F}_d$. By Lemma 4.6.4, we may assume that this U-tree is 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.

(ii) is then immediate from the definition of $\mathcal{F}_d$. $\square$

Note that the nested set constructed in Theorem 4.6.28 does not yet necessarily distinguish any two profiles *efficiently.* However, we can use Theorem 4.6.16 in combination with Theorem 4.6.28 to obtain such a set:

Theorem 4.6.29 (Efficient tree-of-tangles theorem for different order profiles)**.** *Let $\vec{U} = (\vec{U}, \leqslant, ^*, \vee, \wedge, |\cdot|)$ 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 4.6.28. 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_i's 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 an 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 2.3.1. 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 (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 4.6.23 and the fish Lemma 2.3.1 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 (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 2.3.1, 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 4.6.16 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 set 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_i's: 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 $\mathcal{P}$. $\square$

Chapter 5

Abstract separation systems

The last chapter of this thesis is devoted not to tangles or other 'highly connected structures' in some context, but instead to abstract separation systems themselves. A lot of results in this thesis, as well as in tangle theory in general, are formulated in the context of these abstract separation systems, so better understanding them gives a better insight on tangle theory in general.

For example, there are different notions of submodularity of a separation system: a separation system could be structurally submodular in the surrounding universe of separations, or it could be the set S_k for some submodular order function on that universe, and so far no one has shown that these two classes of separation systems are not actually the same, i.e. constructed an explicit example of a structurally submodular separation system for which there does not exist a submodular order function and a natural number k such that this system equals S_k. Such a construction is one of the tasks we perform in Section 5.1. Moreover, we there also show that actually every structurally submodular separation system can be obtained as a disjoint union of systems of bipartitions. These results are published in [38], which is joint work with Jakob Kneip and Maximilian Teegen, however the results presented in Section 5.1 are joint work with Maximilian Teegen only.

Finally, in Section 5.2 we are concerned with a related question: although we have seen in Section 5.1 that not every structurally submodular separation system arises as the S_k for some submodular order function, it is still unclear, how much these two different concepts can differ from each other. More specifically we are concerned with the question of whether the separations in every structurally submodular separation system can be enumerated in such a way that the set of the first i separations, for every integer i, is again a structurally submodular separation system. We construct an example of such a structurally submodular separation system inside a non-distributive universe, for which this is not the case. We also consider this question for separation systems inside distributive universes and show, that there the problem is actually equivalent to a question about systems of sets. Moreover, we solve a variation of that problem, by showing that such an enumeration is indeed possible if one considers a slightly different notion of submodularity. Section 5.2 is joint work with Jakob Kneip and Maximilian Teegen and published in [40].

5.1 The Structure of Submodular Separation Systems

5.1.1 Introduction

My co-authors and I have, in much of our work, relied heavily on structural submodular separation systems. Indeed, separation systems which are structurally 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 [26, 36, 39, 41].

However, so far no one has analysed under which conditions this structural submodularity of a separation system $\vec{S}$ inside a universe $\vec{U}$ actually arises from some submodular order function, that is whether we can find some submodular order function and an integer k such that $\vec{S}$ equals $\vec{S}_k$. If this is the case, we say that the submodularity of $\vec{S}$ in $\vec{U}$ is *order-induced* in $\vec{U}$. This question, under which conditions the submodularity of a separation system is order-induced, is what we approach in Sections 5.1.3 and 5.1.4. In Section 5.1.3 we prove that not every structurally submodular separation systems is order-induced:

Theorem 20. *There exists a separation system $\vec{S}$ which is submodular in a universe $\vec{U}$ of set bipartitions 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 5.1.4 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 5.1.5, 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 (natural) 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 parts' 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 21. *Every separation system $\vec{S}$ which is submodular in some distributive universe $\vec{U}$ of separations is a disjoint union of corner-closed subsystems $\vec{S}_1, \ldots, \vec{S}_n$ of $\vec{S}$ (which are thus also submodular in $\vec{U}$) each of which can be corner-faithfully embedded into a universe of bipartitions.*

Careful analysis of the proof of our decomposition theorem allows us to explicitly specify the subsystems.

The research in this section was inspired, in part, by our search for a solution to the unravelling problem which can be found in Section 5.2.

5.1.2 Preliminaries

Additionally to the terminology given in Chapter 2, we need some essential additional terminology from lattice theory. In this section, we state this terminology and also introduce some definitions specific to this section, most of which are generalizations of definitions made for separation systems and universes of separations to posets and lattices, respectively.

Lattice theory

A *sublattice* L' of a lattice L is a subset of L which is closed under pairwise joins and meets in L.

Recall that an important example of a lattice is the *subset lattice* of a finite set V which consists of all subsets of V, ordered by inclusion. In fact, all finite distributive lattices can be represented as such 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 5.1.1 (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.* [54]

Submodularity and additional terminology

Moreover, we need some terminology specific to this section, mostly generalized versions of established terminology.

A separation system $\vec{S}$ is a *subsystem* of another separation system $\vec{S'}$ if $\vec{S} \subseteq \vec{S'}$ and the involution on $\vec{S}$ is the restriction of the involution on $\vec{S'}$. In particular, $\vec{S}$ is a subset of $\vec{S'}$ which is closed under the involution on $\vec{S'}$. If a subsystem $\vec{S}$ of a universe $\vec{U}$ is closed under joins and meets in $\vec{U}$, we say that $\vec{S}$ (together with the restrictions of $\vee$ and $\wedge$) is a *subuniverse* of $\vec{U}$. For example, the bipartition universe $\mathcal{B}(V)$ on a set V is a subuniverse of the universe of set separations of V.

[21] considers submodular order functions for universes of separations, we will need the more general notion of such a function for arbitrary lattices. Given a lattice L, a function $f\colon L \to \mathbb{R}_0^+$ is called *submodular* if

$$f(a \vee b) + f(a \wedge b) \leqslant f(a) + f(b)$$

for all $a, b \in L$. Note that we here, unlike for separation systems, allow our function to take values in $\mathbb{R}_0^+$ instead of $\mathbb{N}$. However, as already mentioned in the introduction, for finite separation systems we could have also defined an order function to take values in $\mathbb{R}_0^+$ instead of $\mathbb{N}$, due to Lemma 2.4.1.

We say that a subset P is *order-induced submodular* in a lattice L, or that *the submodularity of P in L is order-induced in L* if there exists some submodular function f on L such that $P = \{a \in L : f(a) < k\}$ for some k. In this case, we also say that f *induces the submodularity of P in L* and that f *and k induce the submodularity of P in L*.

Similarly, we say that a subsystem $\vec{S}$ of a universe $\vec{U}$ is *order-induced submodular* in $\vec{U}$ if there exists a submodular order function f on $\vec{U}$ (which takes values in $\mathbb{N}$) and some $k \in \mathbb{N}$ which induce the submodularity of $\vec{S}$ in $\vec{U}$, i.e. such that $\vec{S} = \vec{S}_k$.

Given a lattice L and some subset P of L we say that P is *submodular in L* if, for all $a, b \in P$, the following holds: the supremum of a and b (taken in L) is contained in P', or[1] the infimum of a and b (taken in L) is contained in P'.

5.1.3 Structural submodularity which is not order-induced

In this section we deal with the question of 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 P inside a lattice L is not order-induced is to find a sequence $a_1, a_2, \dots a_n$ of elements of L so that every submodular function f on L for which P is the set of all elements of L of order less than k would need to satisfy $f(a_1) < f(a_2) < \dots f(a_n) < f(a_1)$.

[1]Note that this 'or' is not exclusive.

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 5.1.2. *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 $f(a \wedge c) > f(c)$. Moreover, 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 5.1.3. *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 5.1.4. *Let $P \subseteq L$ be submodular in L, then the dependency digraph of P contains no directed cycle of length two.*

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. □

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. 5.1; a formal description follows.

Consider the universe $\vec{U} = \mathcal{B}(V)$ 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.

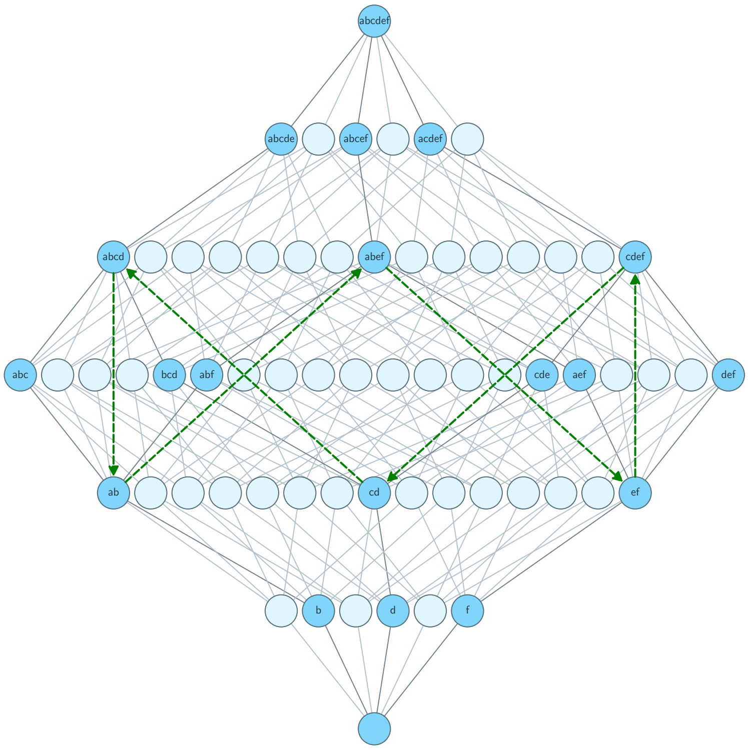

Figure 5.1: The Hasse diagram of $\vec{U}$ from Theorem 20. For readability, only points in $\vec{S}$ are labelled and only one side of each bipartition is denoted.

This example proves the following theorem, as every separation system which is order-induced submodular in a universe $\vec{U}$ is also a poset which is order-induced submodular in the lattice $\vec{U}$:

Theorem 20. *There exists a separation system $\vec{S}$ which is submodular in a universe $\vec{U}$ of set bipartitions 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. 5.2 and indicate the partial lattice inside this lattice as well as the cycle in the 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 5.1.5. Does there exist 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?

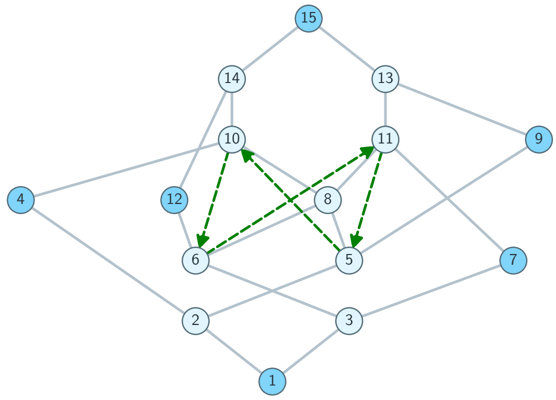

Figure 5.2: The dark blue elements form a partial lattice, which does not contain a cycle in the inner dependency digraph, however, the green dashed edges form a cycle in the outer dependency digraph.

We can ask the same question for a submodular partial lattice:

Question 5.1.6. Does there exist 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 first that a positive answer to Question 5.1.5 implies a positive answer to Question 5.1.6: 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 $k \in \mathbb{R}_0^+$ and $f_l : \vec{U} \to \mathbb{R}_0^+$ would be a submodular function witnessing that $\vec{S}$ is order-induced submodular as a partial lattice, then we could consider the function f given by $f(\vec{s}) = f_l(\vec{s}) + f_l(\overleftarrow{s})$ for every $\vec{s} \in \vec{S}$, which would then, after applying Lemma 2.4.1 result in a submodular order function for $\vec{U}$ as a universe, which witness that the submodularity of $\vec{S}$ in $\vec{U}$ is order-induced.

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 $\vec{U}$.[2] Moreover, $\vec{S}$ is not order-induced submodular, since we can restrict any witnessing submodular

[2]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 $\overleftarrow{s} \leqslant \vec{s}$.

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.

5.1.4 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 in how far 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 subuniverse $\vec{U'}$ of a universe $\vec{U}$ there exists a submodular order function f and some k, such that $\vec{U'} = f^{-1}([0, k])$.

It suffices to first consider these problems for submodular functions on lattices, rather than submodular order functions on universes of separations: due to Lemma 2.4.1 it is enough to find corresponding functions with values in $\mathbb{R}$ instead of $\mathbb{N}$. Now if $f' \colon \vec{U'} \to \mathbb{R}_0^+$ is a submodular function on $\vec{U'} \supseteq \vec{U}$ which agrees on $\vec{U}$ with some submodular *symmetric* function $f \colon \vec{U} \to \mathbb{R}_0^+$, then we can define a submodular *symmetric* 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, by Lemma 2.4.1, we find a corresponding order function as desired.

For the first theorem, recall that an *interval* in a lattice L is, for some $x, y \in L$, the 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 5.1.7. *Let L be a lattice and $L' = [x, y] \subseteq L$ an interval in L. Suppose that $f \colon L' \to \mathbb{R}_0^+$ is a submodular function on L' with maximum value k. Then there exists a submodular function $g \colon L \to \mathbb{R}_0^+$ 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* $dl(z)$ recursively as follows: assign $dl(\bot) = 0$ for the bottom element $\bot$ of L. Now

$$dl(z) := \max\{dl(z') + 1 \,:\, z' < z\}$$

for all other $z \in L$. Similarly, for $z \in L$ such that $z \geqslant x$ we define its *up-level* $ul(z)$ recursively: we assign $ul(\top) = 0$ for the top element $\top$ of L. Now

$$ul(z) := \max\{ul(z') + 1 \ : \ z' > z\}.$$

Let l be the maximum possible level (up or down) and let $M = 2 \cdot 2^l \cdot k > k$. We now define g as follows:

$$g(z) = \begin{cases} f(z) & z \in L' \\ M \cdot (2 - 2^{-dl(z)}) & z \in L^{\downarrow} \\ M \cdot (2 - 2^{-ul(z)}) & z \in L^{\uparrow} \\ 4 \cdot M & z \in L^{\leftrightarrow} \end{cases}$$

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 $ul(a) > ul(a \vee b)$, thus

$$\begin{aligned} g(a \vee b) + g(a \wedge b) &\leqslant M \cdot (2 - 2^{-ul(a \vee b)}) + 4 \cdot M \\ &< M \cdot (2 - 2^{-ul(a)}) + 4 \cdot M = g(a) + g(b). \end{aligned}$$

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 $ul(a) \leqslant ul(b)$. By the definition of $L^{\uparrow}$ and ul, we have $a \vee b \in L^{\uparrow}$ and $ul(a \vee b) < ul(a)$. Furthermore, $a \wedge b \in L^{\uparrow} \cup L'$, so in any case $g(a \wedge b) < 2M$. We calculate

$$\begin{aligned} g(a \vee b) + g(a \wedge b) &< M \cdot (2 - 2^{-(ul(a)-1)}) + 2M \\ &= 4M - M(2^{-ul(a)} + 2^{-ul(a)}) \\ &\leqslant 4M - M(2^{-ul(a)} + 2^{-ul(b)}) = g(a) + g(b). \end{aligned}$$

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 $z \in L \smallsetminus L'$, by the definition of M, the function g is as claimed. □

This theorem will also serve as a tool in proving the second theorem, which is the following:

Theorem 5.1.8. *Let L be a distributive lattice and $L' \subseteq L$ a sublattice. Then there exists a submodular function $f\colon L \to \mathbb{R}_0^+$ and a $k \in \mathbb{R}_0^+$ such that $L' = f^{-1}([0,k])$.*

Theorem 5.1.7 allows us to first prove Theorem 5.1.8 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 5.1.7.

Lemma 5.1.9. *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 $f\colon L \to \mathbb{R}_0^+$ such that $L' = f^{-1}(0)$.*

Proof. By the Birkhoff representation theorem (Theorem 5.1.1) 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,

$$f(Y) = \sum_{p \in Y} |X_p \smallsetminus Y|.$$

This function is submodular, since for all $X, Y \in L$ we can calculate as follows

$$\begin{aligned}
&f(X) + f(Y) \\
=&\sum_{p \in Y} |X_p \smallsetminus Y| + \sum_{p \in X} |X_p \smallsetminus X| \\
=&\sum_{p \in X \cap Y} (|X_p \smallsetminus Y| + |X_p \smallsetminus X|) + \sum_{p \in Y \smallsetminus X} |X_p \smallsetminus Y| + \sum_{p \in X \smallsetminus Y} |X_p \smallsetminus X| \\
=&\sum_{p \in X \cap Y} (|X_p \smallsetminus (X \cap Y)| + |X_p \smallsetminus (X \cup Y)|) \\
&+ \sum_{p \in Y \smallsetminus X} |X_p \smallsetminus Y| + \sum_{p \in X \smallsetminus Y} |X_p \smallsetminus X| \\
=&f(X \cap Y) + \sum_{p \in X \cap Y} |X_p \smallsetminus (X \cup Y)| + \sum_{p \in Y \smallsetminus X} |X_p \smallsetminus Y| + \sum_{p \in X \smallsetminus Y} |X_p \smallsetminus X| \\
\geqslant&f(X \cap Y) + \sum_{p \in X \cap Y} |X_p \smallsetminus (X \cup Y)| \\
&+ \sum_{p \in Y \smallsetminus X} |X_p \smallsetminus (X \cup Y)| + \sum_{p \in X \smallsetminus Y} |X_p \smallsetminus (X \cup Y)| \\
=&f(X \cap Y) + \sum_{p \in X \cup Y} |X_p \smallsetminus (X \cup Y)| \\
=&f(X \cap Y) + f(X \cup Y).
\end{aligned}$$

Thus, all that is left to show is that $f(Y) > 0$ for every $Y \in L \smallsetminus L'$. 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 $X_p \not\subseteq Y$. In particular there needs to be some $q \in X_p$ such that $q \notin Y$, which witnesses that $f(Y) > 0$. □

Combining Lemma 5.1.9 and Theorem 5.1.7 results in a proof of Theorem 5.1.8:

Proof of Theorem 5.1.8. Let $\bot$ be the bottom element of L' and let $\top$ be the top element of L'. By Lemma 5.1.9 there is a submodular function $f : L'' \to \mathbb{R}_0^+$ on $L'' = [\bot, \top] \supseteq L'$ such that $f^{-1}(0) = L'$. Using this f as input in Theorem 5.1.7 results in the desired submodular function on L. □

From Theorem 5.1.7 and Theorem 5.1.8 we now immediately obtain the same results for subuniverses, in the way discussed above:

Theorem 5.1.10. *Given a distributive universe $\vec{U}$ of separations and a subuniverse $\vec{U'} \subseteq \vec{U}$, there is a submodular order function $f\colon \vec{U} \to \mathbb{N}$ and a $k \in \mathbb{N}$ such that $\vec{U'} = f^{-1}([0, k])$.*

Proof. We apply Theorem 5.1.8 to $\vec{U'}$ as a sublattice of $\vec{U}$ to obtain a submodular function f' and $k' \in \mathbb{R}_0^+$ with $\vec{U'} = f'^{-1}([0, k'])$. We now define a symmetric order function f on $\vec{U}$ with $f(\vec{s}) := f'(\vec{s}) + f'(\overleftarrow{s})$. Applying Lemma 2.4.1 then results in the desired function, as can be seen by applying that lemma with $k := 2k'$, since $\vec{U'} = f^{-1}([0, k])$. □

Theorem 5.1.11. *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 $f \colon \vec{U'} \to \mathbb{N}$ is a submodular order function on $\vec{U'}$ with maximum value k. Then there exists a submodular order function $g \colon \vec{U} \to \mathbb{N}$ 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 5.1.7 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 $\vec{U'}$. We notice from the proof of Theorem 5.1.7 that $g'(\vec{z})$ is an even natural number for every $\vec{z} \notin \vec{U'}$. Now this function g' 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 and takes it values in $\mathbb{N}$. Since g' takes values larger than k outside of $\vec{U'}$, so does g. □

5.1.5 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 pigeonhole principle, at least one of these subsystems will have a size of at least $\frac{|S|}{3}$. This observation vaguely links the question of submodular decompositions to the unravelling problem (see Section 5.2): 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 $\{\vec{s}, \overleftarrow{s}, \vec{s} \vee \overleftarrow{s}, \vec{s} \wedge \overleftarrow{s}\}$.

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 subsystem 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 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.

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.

Theorem 5.1.12. *Given a universe $\vec{U} = \mathcal{B}(V)$ 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. Since $|S| \geqslant 3$, 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, these 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) is 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, 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. □

Birkhoff's theorem for distributive universes and decompositions in distributive universes

To lift Theorem 5.1.12 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 5.1.4, use the Birkhoff representation theorem for distributive lattices:

Theorem 5.1.1 (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 5.1.13 (Birkhoff representation of universes of separations)**.**
For every involution poset[3] $(P, \leqslant, ')$, 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 U and $(\mathcal{O}(P), {}^*)$.

Proof of Theorem 5.1.13. 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 $P \coloneqq \mathcal{J}(\vec{U})$ so that $\vec{U}$ is isomorphic to $\mathcal{O}(P)$.

Theorem 5.1.1 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-closed 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:

$$\textit{For all } X \in \mathcal{O}(P) \textit{ we have that } |X^*| = |P| - |X|. \tag{$\dagger$}$$

We prove Eq. (†) by contradiction. So assume that X is an inclusion-wise minimal down-closed subset of P for which Eq. (†) 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 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 Eq. (†).

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 Eq. (†). 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

[3] Recall that involution posets are the same as separation systems. However, to emphasize 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.

$\lceil x \rceil^* = P \smallsetminus \lfloor x' \rfloor$, we have that $(X - x)^* \subseteq \lceil x \rceil^*$ and thus $(X - x) \supseteq \lceil x \rceil$, which is absurd.

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

$$((\lceil x \rceil - x)^* - x')^* \smallsetminus (\lceil x \rceil - x)^{**} = \lceil x \rceil \smallsetminus (\lceil x \rceil - x),$$

so x'' is x.

Let us show that we have $X^* = P \smallsetminus \{x' : x \in X\}$ 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 $(X - x)^* = P \smallsetminus \{y' : y \in (X - x)\}$. By the earlier observation, the single element in $(X - x)^* \smallsetminus X^*$ is precisely x', giving $X^* = P \smallsetminus \{y' : y \in X\}$ 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 $\lceil x \rceil^* \subseteq P \smallsetminus \lfloor x' \rfloor$. By applying * to both sides and using the above paragraph we get that $\lceil x \rceil \supseteq P \smallsetminus \{y' : y \in P \smallsetminus \lfloor x' \rfloor\}$. 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 $y' \geqslant x'$. □

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 5.1.14. *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 is by induction on $|\vec{U}|$.

By applying Theorem 5.1.13 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

$$\begin{aligned} \vec{S}_p &\coloneqq \{X \in \vec{S} : p \in X,\, p' \notin X\}, \\ \vec{S}_{p'} &\coloneqq \{X \in \vec{S} : p \notin X,\, p' \in X\}, \\ \vec{S}_{p,p'} &\coloneqq \vec{S} \smallsetminus (\vec{S}_p \cup \vec{S}_{p'}). \end{aligned}$$

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 $|\vec{U'}| < |\vec{U}|$ 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 5.1.12 gives the desired decomposition. □

Observe that in $(\mathcal{O}(P), {}^*)$ we have $X \wedge X^* = \{p \in P \,:\, p \in X, p' \notin X\}$. Hence, by 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 21. *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 $\vec{U}$.

5.2 The Unravelling Problem

5.2.1 Introduction

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 22 (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 $|\mathcal{X}_i \smallsetminus \mathcal{X}_{i-1}| = 1$ 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 of whether every woven set has an unravelling arose naturally in our study of structurally submodular separation systems, which we will make precise in Section 5.2.3.

In this section we analyse the unravelling problem. We prove affirmative versions in two important cases, which come from the original context of 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 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.

We start in Section 5.2.2 with the additional definitions required for this section and show that the unravelling problem has an equivalent formulation in terms of distributive lattices. We 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.

In Section 5.2.3 we relate the unravelling problem to abstract separation systems, and explain how it naturally arises in that context. Throughout the rest of this section, we will come back to the context of abstract separation systems, to discuss how our results apply there.

In Section 5.2.4 we give a partial solution to the unravelling problem, by showing that those woven sets, which arise as an S_k for a submodular order function, 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 $f : 2^V \to \mathbb{R}$ and $k \in \mathbb{R}$, let us say that f *induces* $\mathcal{X}$.

Theorem 5.2.1. *If $\mathcal{X} \subseteq 2^V$ is induced by a submodular function on 2^V, then $\mathcal{X}$ can be unravelled.*

In Section 5.2.5 we introduce a possible generalization 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 22.

We conclude in Section 5.2.6 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 infimum 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:

Theorem 23. *Every woven poset can be unravelled.*

Wovenness in posets corresponds to the most general notion of submodularity for separation systems, which is also discussed in [38].

5.2.2 Preliminaries

Given a poset P and $a, b \in P$ we say that a is an *upper (lower) cover* of b if $a > b$ ($a < b$) and there does not exist any $c \in P$ such that $a > c > b$ ($a < c < b$).

Using this, we can formulate a lattice-theoretical variation of the unravelling problem. If L is a lattice, we say that $P \subseteq L$ is *woven* in L if for any $p, q \in P$ we have $p \vee q \in P$ or $p \wedge q \in P$.

We can now generalize Problem 22 as follows:

Problem 5.2.2. 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 5.2.2 is equivalent to Problem 22 by Birkhoff's representation Theorem 5.1.1 (see [16]), 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 5.2.2: in Section 5.2.5 we shall construct a (non-distributive) counterexample for Problem 5.2.2.

Perhaps surprisingly, it is easy to establish a kind of converse to Problem 5.2.2: 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 5.2.3. *If L is a lattice and $P \subsetneq L$ a proper woven subset of L, then there is 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 $L \setminus P$. □

In terms of woven sets in the sense of Problem 22, this statement directly implies the following:

Corollary 5.2.4. *If V is a finite set and $\mathcal{X} \subsetneq 2^V$ woven, then there is an $X \subseteq V$ such that $X \notin \mathcal{X}$ and such that $\mathcal{X} + X$ is again woven.*

5.2.3 Abstract separation systems and submodularity

Our original motivation for considering the unravelling problem originates in structural submodular separation systems: given a structural 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 structurally submodular in $\vec{U}$. Formally, given a structurally submodular separation system $\vec{S}$ inside a universe $\vec{U}$ of separations, we are interested if the following property holds:

Property 5.2.5. 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 structurally submodular separation systems, 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 structurally submodular separation systems 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 structurally submodular separation systems via induction. For example, it is possible to obtain a short proof of a tree-of-tangles theorem for structurally submodular separation systems via this unravelling sequence [60, Section 4.1.8].

This question, whether Property 5.2.5 holds for every structurally submodular separation system, is now closely related to Problem 5.2.2. In fact, if we could unravel every structurally submodular separation system, we could answer Problem 5.2.2 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 5.2.5 to turn our counterexample to Problem 5.2.2 into an example of a structurally submodular separation system inside a non-distributive lattice which cannot be unravelled.

Also, the converse of Problem 5.2.2 established in Proposition 5.2.3 directly translates to a similar statement about structurally submodular separation systems inside a universe of separations.

Corollary 5.2.6. *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 5.2.3, the separation system $S' := S + r$ is again submodular in U. □

5.2.4 Unravelling order-induced sets

In this section we show that for a subclass of the woven subsets of a lattice we indeed have unravellings.

Recall that a subset P of a lattice L is *order-induced submodular* if there exists a submodular function $f : L \to \mathbb{R}_0^+$ and a real number k such that $P = \{p \in L : f(p) < k\}$. Here, f being submodular means 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 exist woven sets which are not order-induced submodular, for example given by the structurally submodular separation system constructed in Section 5.1.3.

We will see in what follows that for order-induced submodular posets P 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, so 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 subset of a lattice can be unravelled.

Theorem 5.2.7. *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 5.2.8. *If $P = \{p \in L : f(p) < k\}$ and $p \in P$ maximizes $f(p)$ 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 $f(p) \geqslant f(q), f(r)$. Thus, if one of $f(q \vee r)$ and $f(q \wedge r)$ equals $f(p)$, the other also needs to lie in P. Thus, $P - p$ is indeed woven in L. $\square$

Unfortunately we cannot rely solely on Lemma 5.2.8 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 5.2.8, 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 [68] to construct certain tree-decompositions. For this we show the following:

Theorem 5.2.9. *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 $I(p_1), I(p_2), I(q_1), I(q_2)$ 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 $q_1 < p_2$. Since $p_2 \in I(p_2)$ this then implies $p_2 < q_2$. Since $q_2 \in I(q_2)$ 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 5.2.10. *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 5.2.9 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 5.2.9 guarantees that indeed $\gamma(r) \neq \gamma(s)$ for all $r \neq s \in U$. □

We can now establish Theorem 5.2.7.

Proof of Theorem 5.2.7. Let L be a lattice with a submodular function f. Let $P = \{p \in L \,:\, f(p) < k\}$ for some $k \in \mathbb{R}_0^+$. Let ρ be the submodular function on L from Theorem 5.2.9. Let ϵ be the minimal difference between two distinct values of f, that is $|f(p) - f(q)| \geqslant \epsilon$ or $f(p) = f(q)$ for any two $p, q \in L$. Since L is finite, $\epsilon > 0$. Pick a positive constant $c \in \mathbb{R}^+$ so that $c \cdot \rho(p) < \epsilon$ for all $p \in L$. We define a new function $g\colon L \to \mathbb{R}_0^+$ 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 of P as $p_1, \dots, p_n$ in such a way that $g(p_1) < g(p_2) < \dots < g(p_n)$. 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 an unravelling for P is given by $P = P_n \supseteq \dots \supseteq P_0 = \emptyset$. □

Theorem 5.2.7 allows us to give a class of sets $\mathcal{X} \subseteq 2^V$ for which we can answer Problem 22 positively. We say that a function $f : 2^V \to \mathbb{R}$ is *submodular* if $f(X) + f(Y) \geqslant f(X \cup Y) + f(X \cap Y)$ for all $X, Y \in 2^V$ and obtain the following theorem as a corollary:

Theorem 5.2.1. *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 $f(X) \geqslant 0$ for all $X \subseteq V$. Applying Theorem 5.2.7 to the subset-lattice 2^X together with its subset $\mathcal{X}$ results in the desired unravelling sequence $\emptyset = \mathcal{X}_0 \subseteq \mathcal{X}_1 \subseteq \dots \subseteq \mathcal{X}_n = \mathcal{X}$. □

Moreover, Theorem 5.2.7 also allows us to show that separation systems $\vec{S}_k$ inside a universe of separations with a submodular order function can be unravelled.

Corollary 5.2.11. *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 5.2.7, using the function γ from Corollary 5.2.10 instead of the function ρ from Theorem 5.2.9. $\square$

5.2.5 A woven subset of a lattice which cannot be unravelled

In this section we are going to construct a counterexample to Problem 5.2.2 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 5.2.12 ([61])**.** *There exists a* 4*-regular graph G with girth at least* 11.

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 5.2.13. *If G is a graph, $W = v_1 v_2 \dots v_n v_1$ a closed walk in G such that there exists a 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. $\square$

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 11. The ground set of our lattice L consists of a top element $\top$, a bottom element $\bot$, and 4 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 12. 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 $\bot \leqslant v$ and $v \leqslant \top$ for every $v \in L$. It is easy to see that this $\leqslant$ is indeed a partial order.

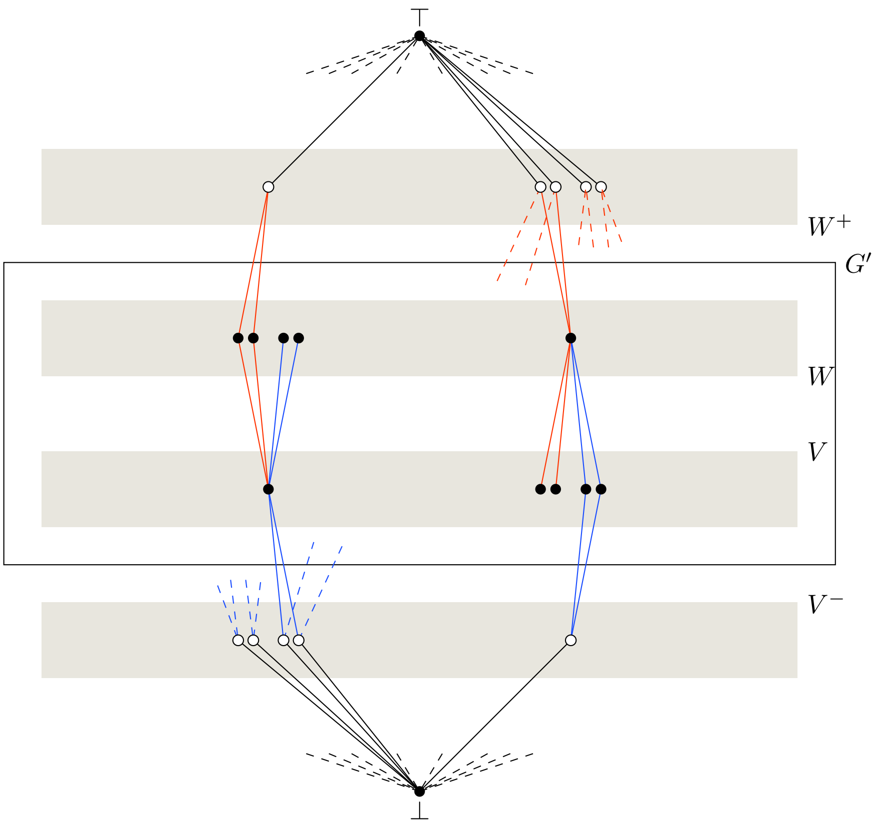

Figure 5.3: The Hasse diagram of L. The points in P are denoted by black dots, the points outside of P are white.

We claim that $(L, \leqslant)$ is a lattice, that $P = V \cup W \cup \{\top, \bot\} \subseteq L$ 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:

Lemma 5.2.14. *If $v_1, v_2 \in V$, then there is a supremum and an infimum of v_1, v_2 in L. Moreover, if $v_1 \wedge v_2 \neq \bot$ 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* $w_1 \vee w_2 \neq \top$*, 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 $x = \top$, then this is clear and $x \in W \cup V \cup V^- \cup \{\bot\}$ 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. Since $v_1 \leqslant x$ we may suppose without loss of generality that $v_1 \leqslant w_1$. Now if $v_2 \leqslant w_2$, then $w v_1 w_1 v_s w_2 v_2 w$ contains a cycle of length at most 6 in G' by Lemma 5.2.13, 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 $\top$, 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 $w_{21}, w_{22} \leqslant w_2^+$.

Since $v_1 \leqslant w_1^+, w_2^+$, we may also 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 4. In any case, we consider the closed walk $v_1 w_{11} v_{w_1^+} w_{12} v_2 w_{22} v_{w_2^+} w_{21} v_1$. Since $v_{w_1^+} \neq v_{w_2^+}$, 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 5.2.13 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 $\bot$. 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 $v_1 \wedge v_2 \neq \bot$, then there is a point $w \in W$ such that both, $v_1 w$ and $v_2 w$ 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 5.2.14 to show the existence of suprema and infima between $v \in V$ and $w \in W$:

Lemma 5.2.15. *If $v \in V$ and $w \in W$, then there is a supremum and an infimum of v and w in L. Moreover, if $v \wedge w \neq \bot$, then $v \vee w = \top$ or $v \leqslant w$.*

Proof. If $v \leqslant w$, then the statement is obvious, so suppose that $v \nleqslant w$.

By Lemma 5.2.14, every point $w_i \in N_{G'}(v)$ has a supremum with w which is either $\top$ 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 5.2.14, $w_1 \wedge w$ 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 occur at the same time.

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 $v \nleqslant w$ 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 5.2.13 to our walk to show the existence of a cycle of length at most 6 in G', which is a contradiction. □

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 5.2.16. *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 $v^- \wedge x = \bot$.

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 $x \vee v^- = \top$.

If $x \in W$, then $x \vee v_1$ and $x \vee v_2$ exist by Lemma 5.2.15 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 $v_1 \nleqslant x \vee v_2$ as well as $v_2 \nleqslant x \vee v_1$. 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 $v_1 \nleqslant x$ and $v_2 \nleqslant x$ 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 5.2.15, 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 $x w_2 v_{x_2} w_1 v_2 w_{v^-} v_1 (x \vee v_1) x$ contains a cycle of length at most 8 in G' by Lemma 5.2.13, 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 $x \vee v_1$ and let v_{x_2} be the point in V corresponding to $x \vee v_2$. Further, 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, $x w_2 v_{x_1} w_1 v_1 w_{v^-} v_2 w_3 v_{x_2} w_4 x$ contains a cycle in G' of length at most 10 by Lemma 5.2.13.

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 5.2.15, 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 5.2.17. *L is a lattice and $P = V \cup W \cup \{\top, \bot\} \subseteq L$ is woven in L such that $P - p$ is not woven in L for any $p \in P$.*

Proof. By the Lemmas 5.2.14 to 5.2.16 L is indeed a lattice. To see that P is woven in L observe that by Lemma 5.2.14, Lemma 5.2.15 and the fact that $\top$ and $\bot$ 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 $p = \bot$ we note that there are $v_1, v_2 \in V$ such that $v_1 \vee v_2 \in W^+$ which implies that $v_1 \wedge v_2 = \bot$ and shows that $P - \bot$ is not woven in L. Similarly, $P - \top$ 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 5.2.18. *There exists a universe $\vec{U}$ of separations and a submodular subsystem $\vec{S} \subseteq \vec{U}$ such that $\vec{S} - \{\vec{s}, \overleftarrow{s}\}$ 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 $\top$ in L (the top of L) with the copy of $\bot$ in L' (the top of L') and the copy of $\bot$ in L with the copy of

$\top$ 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 $\vec{S} - \{\vec{s}, \overleftarrow{s}\}$ is again submodular in $\vec{U}$. □

Note that neither our lattice L nor the constructed universe $\vec{U}$ of separations are distributive.

5.2.6 Woven posets

Instead of asking in Problem 5.2.2 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 an $r \in P$ 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 an $r \in P$ 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* [64] 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 lattices. 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 5.2.19. *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, 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. □

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 5.2.20. *Let P be a woven poset and M its set of maximal elements. Then every non-empty subset $M' \subseteq M$ has an infimum* $\inf M'$ *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 $M'' \coloneqq M' - m$. 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. □

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 5.2.21. *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* $\inf M'$ *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 lies 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 5.2.22. *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 5.2.21, given any $M' \subseteq M$, there exists at most one element in $d(M')$. Moreover, by Lemma 5.2.20 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 5.2.5) we can indeed unravel every such poset: given some woven poset P, by Lemma 5.2.22, P contains an element p which has only one upper or lower cover, and, by Lemma 5.2.19, $P - p$ is again woven. Thus, we obtain the following theorem:

Theorem 23. *Every woven poset can be unravelled.*

Again we can translate this result to abstract separation systems.

We say that a separation system $\vec{S}$, on its own, not in the context of a surrounding universe $\vec{U}$ of separations, is *submodular* if there exists, for any two separations $\vec{s}, \vec{t} \in \vec{S}$ a supremum or an infimum in $\vec{S}$, i.e. – as for woven posets – we require that there either is a smallest separation $\vec{r}$ such that $\vec{s}, \vec{t} \leqslant \vec{r}$ or there is a largest separation $\vec{r}$ such that $\vec{s}, \vec{t} \geqslant \vec{r}$. These submodular separation systems are also considered in [38], where it was shown that one can find, for each such system $\vec{S}$, a universe $\vec{U}$ of separations in which we can embed $\vec{S}$ so that the joins and meets in $\vec{S}$ are preserved.

We now obtain the following corollary for this type of separation system:

Theorem 5.2.23. *Let $\vec{S}$ be a submodular separation system. Then there exists an $\vec{s} \in \vec{S}$ such that $\vec{S} \smallsetminus \{\vec{s}, \overleftarrow{s}\}$ 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, $\inf M \geqslant \overleftarrow{t}$ for all $\vec{t} \in M$ and thus $\inf M \geqslant \sup M^* = (\inf M)^*$. 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 5.2.21.

We note that $\vec{x} \neq (\inf M')^*$ as otherwise $\vec{x} \leqslant (\inf M)^* \leqslant \inf M$, contradicting the fact that $\vec{x} \in d(M')$. But this implies that $\vec{S} - \vec{x}$ is a woven poset by Lemma 5.2.19. 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 5.2.19, 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 subsets M' of M. This implies that every element $\vec{s} \in \vec{S}$ is nested with $\inf M$: if $\vec{s} \in d(M)$, then $\vec{s} \leqslant \inf M$ and if $\vec{s} \in d(M')$ for a proper subset M' of M, then $\vec{s} = \inf M' \geqslant \inf M$. Now suppose that $|M| \geqslant 2$. Then there is an $\vec{m} \in M$ such that $\overleftarrow{m} \neq \inf M$. We claim that $\vec{S} \smallsetminus \{\vec{m}, \overleftarrow{m}\}$ 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). Since $\vec{x}$ and $\vec{y}$ are nested with $\inf M$ this implies that $\vec{x}, \vec{y} \geqslant \inf M$ as $\vec{x} \leqslant \inf M$ would imply that $\vec{x} \vee \vec{y} = \vec{y}$ or $\vec{x} \vee \vec{y} \leqslant \inf M$. Thus, $\vec{x} = \inf M'$ and $\vec{y} = \inf M''$ for subsets M', M'' of M, say. Thus, $\inf(M' \cup M'')$, which exists by Lemma 5.2.20, is also the infimum of $\vec{x}$ and $\vec{y}$. Moreover, since $\overleftarrow{m} \neq \inf M$ and $\overleftarrow{m}$ is a minimal element of $\vec{S}$ and $\inf(M' \cup M'') \geqslant \inf M$ we have that $\inf(M' \cup M'') \neq \overleftarrow{m}$ and thus there is a corner of $\vec{x}$ and $\vec{y}$ in $\vec{S} \smallsetminus \{\overrightarrow{m}, \overleftarrow{m}\}$.

It remains the case that $|M| = 1$, say $M = \{\overrightarrow{m}\}$. In this case however, we have that $\vec{s} \leqslant \overrightarrow{m}$ for every $\vec{s} \in \vec{S}$. If $\vec{S} = \{\overrightarrow{m}, \overleftarrow{m}\}$ the statement is trivial, so let $\vec{s} \in \vec{S} - \overrightarrow{m}$ 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 $\overrightarrow{m}$ is the unique upper cover of $\vec{s}$. Thus, $\vec{S} - \vec{s}$ is a woven poset by Lemma 5.2.19. 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 5.2.19, and thus $S - s$ is a submodular separation system. □

The Dedekind-MacNeille completion of posets [16] allows us to embed every woven poset into a lattice so that the poset is woven in this lattice. It is shown in [38] that this technique 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 5.2.2.

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 $r, s \in P \smallsetminus \{p\}$, 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 subposet of L.)

Bibliography

[1] A. Berry, R. Pogorelcnik, and G. Simonet, *An introduction to clique minimal separator decomposition*, Algorithms **3** (2010), no. 2, 197–215, doi:10.3390/a3020197.

[2] N. Bowler, Chr. Elbracht, J. Erde, J. P. Gollin, K. Heuer, M. Pitz, and M. Teegen, *Ubiquity in graphs II: Ubiquity of graphs with nowhere-linear end structure* (2018), available at arXiv:1809.00602.

[3] N. Bowler and J. Kneip, *Separations of sets*, Order **34** (2020), 411–425, doi:10.1007/s11083-019-09512-4.

[4] C. Bürger and J. Kurkofka, *Duality theorems for stars and combs II: Dominating stars and dominated combs* (2020), available at arXiv:2004.00593.

[5] ———, *Duality theorems for stars and combs III: Undominated combs* (2020), available at arXiv:2004.00592.

[6] ———, *Duality theorems for stars and combs IV: Undominating stars* (2020), available at arXiv:2004.00591.

[7] ———, *Duality theorems for stars and combs I: Arbitrary stars and combs*, Journal of Graph Theory, posted on 2021, doi:10.1002/jgt.22706.

[8] J. Carmesin, *A short proof that every finite graph has a tree-decomposition displaying its tangles*, European Journal of Combinatorics **58** (2016), 61–65, doi:10.1016/j.ejc.2016.04.007.

[9] ———, *All graphs have tree-decompositions displaying their topological ends*, Combinatorica **39** (2019), 545–596, doi:10.1007/s00493-018-3572-0.

[10] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark, *k-blocks: A connectivity invariant for graphs*, SIAM Journal on Discrete Mathematics **28** (2014), no. 4, 1876–1891, doi:10.1137/130923646.

[11] ———, *Canonical tree-decompositions of finite graphs I. Existence and algorithms*, Journal of Combinatorial Theory, Series B **116** (2016), 1–24, doi:10.1016/j.jctb.2014.04.001.

[12] ———, *Canonical tree-decompositions of finite graphs II. Essential parts*, Journal of Combinatorial Theory, Series B **118** (2016), 268–283, doi:10.1016/j.jctb.2014.12.009.

[13] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein, *Connectivity and tree structure in finite graphs*, Combinatorica **34** (2014), 1–35, doi:10.1007/s00493-014-2898-5.

[14] J. Carmesin, M. Hamann, and B. Miraftab, *Canonical trees of tree-decompositions* (2020), available at arXiv:2002.12030.

[15] D. Coudert and G. Ducoffe, *Clique-decomposition revisited* (2016). available at hal-01266147.

[16] B. A. Davey and H. A. Priestley, *Introduction to lattices and order*, 2nd ed., Cambridge University Press, 2002.

[17] W. Dicks and M. J. Dunwoody, *Groups acting on graphs*, Cambridge Studies in Advanced Mathematics, Cambridge University Press, 1989.

[18] Y. Dieng, *décomposition arborescente des graphes planaires et routage compact*, Doctorate thesis, Université Bordeaux 1, 2009. available at `http://ori-oai.u-bordeaux1.fr/pdf/2009/DIENG_YOUSSOU_2009.pdf`.

[19] R. Diestel, *Ends and tangles*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **87** (2017), 223–244, doi:10.1007/s12188-016-0163-0.

[20] ———, *Graph theory*, 5th ed., Graduate Texts in Mathematics, vol. 173, Springer, 2017.

[21] ———, *Abstract separation systems*, Order **35** (2018), 157–170, doi:10.1007/s11083-017-9424-5.

[22] ———, *Tree sets*, Order **35** (2018), 171–192, doi:10.1007/s11083-017-9425-4.

[23] R. Diestel, Ph. Eberenz, and J. Erde, *Duality theorem for blocks and tangles in graphs*, SIAM Journal on Discrete Mathematics **31** (2017), no. 3, 1514–1528, doi:10.1137/16M1077763.

[24] R. Diestel, Chr. Elbracht, J. Erde, and M. Teegen, *Duality and tangles of set partitions*, In preparation.

[25] R. Diestel, Chr. Elbracht, and R. Jacobs, *Deciders for tangles of set partitions*, In preparation.

[26] R. Diestel, J. Erde, and D. Weißauer, *Structural submodularity and tangles in abstract separation systems*, Journal of Combinatorial Theory, Series A **167C** (2019), 155–180, doi:10.1016/j.jcta.2019.05.001.

[27] R. Diestel, F. Hundertmark, and S. Lemanczyk, *Profiles of separations: in graphs, matroids, and beyond*, Combinatorica **39** (2019), 37–75, doi:10.1007/s00493-017-3595-y.

[28] R. Diestel, T. R. Jensen, K. Yu. Gorbunov, and C. Thomassen, *Highly connected sets and the excluded grid theorem*, Journal of Combinatorial Theory, Series B **75** (1999), 61–73.

[29] R. Diestel and S.-i. Oum, *Tangle-tree duality in graphs, matroids and beyond*, Combinatorica **39** (2019), 879–910, doi:10.1007/s00493-019-3798-5.

[30] ———, *Tangle-tree duality in abstract separation systems*, Advances in Mathematics **377** (2021), 107470, doi:10.1016/j.aim.2020.107470.

[31] R. Diestel and G. Whittle, *Tangles and the Mona Lisa* (2016), available at arXiv:1603.06652.

[32] G. Ding, *Graphs without large $K_{2,n}$-minors* (2017), available at arXiv:1702.01355.

[33] M. J. Dunwoody and B. Krön, *Vertex cuts*, Journal of Graph Theory **80** (2015), no. 2, 136–171, doi:10.1002/jgt.21844.

[34] Ph. Eberenz, *Characteristics of profiles*, Master's thesis, Universität Hamburg, 2015.

[35] Chr. Elbracht, *Tangles determined by majority vote*, Master's thesis, Universität Hamburg, 2020.

[36] Chr. Elbracht and J. Kneip, *A canonical tree-of-tangles theorem for submodular separation systems*, Combinatorial theory (to appear), available at arXiv:2009.02091.

[37] Chr. Elbracht, J. Kneip, and M. Teegen, *Tangles are decided by weighted vertex sets*, Advances in Combinatorics **2020:9** (2020), doi:10.19086/aic.13691.

[38] ———, *The structure of submodular separation systems* (2021), available at arXiv:2103.13162.

[39] ———, *Trees of tangles in abstract separation systems*, Journal of Combinatorial Theory, Series A **180** (2021), 105425, doi:10.1016/j.jcta.2021.105425.

[40] ———, *The unravelling problem* (2021), available at arXiv:2103.12565.

[41] ———, *Obtaining trees of tangles from tangle-tree duality*, Journal of Combinatorics (to appear), available at arXiv:2011.09758.

[42] ———, *Trees of tangles in infinite abstract separation systems*, Mathematical Proceedings of the Cambridge Philosophical Society (to appear), available at arXiv:2005.12122.

[43] Chr. Elbracht, J. Kurkofka, and M. Teegen, *Edge-connectivity and tree-structure in finite and infinite graphs* (2020), available at arXiv:2012.07651.

[44] M. N. Ellingham, E. A. Marshall, K. Ozeki, and S. Tsuchiya, *A characterization of $K_{2,4}$-minor-free graphs*, SIAM Journal on Discrete Mathematics **30(2)** (2016), 955–975, doi:10.1137/140986517.

[45] A.-K. Elm and J. Kurkofka, *A tree-of-tangles theorem for infinite tangles* (2020), available at arXiv:2003.02535.

[46] J. Erde, *Refining a tree-decomposition which distinguishes tangles*, SIAM Journal on Discrete Mathematics **31** (2017), no. 3, 1529–1551, doi:10.1137/16M1059539.

[47] J. Geelen, *Personal communication*, Oberwolfach, 2016.

[48] J. Geelen, B. Gerards, N. Robertson, and G. Whittle, *Obstructions to branch-decomposition of matroids*, Journal of Combinatorial Theory, Series B **96** (2006), no. 4, 560–570, doi:10.1016/j.jctb.2005.11.001.

[49] J. Geelen, B. Gerards, and G. Whittle, *Tangles, tree-decompositions and grids in matroids*, Journal of Combinatorial Theory, Series B **99** (2009), 657–667, doi:10.1016/j.jctb.2007.10.008.

[50] J. Geelen and B. Joeris, *A generalization of the grid theorem*, available at arXiv:1609.09098.

[51] A. C. Giannopoulou, K.-i. Kawarabayashi, S. Kreutzer, and O-j. Kwon, *The canonical directed tree decomposition and its applications to the directed disjoint paths problem* (2020), available at arXiv:2009.13184.

[52] J. P. Gollin and J. Kneip, *Representations of infinite tree sets*, Order **38** (2021), 79–96, doi:10.1007/s11083-020-09529-0.

[53] M. Grohe, *Quasi-4-connected components*, 43rd international colloquium on automata, languages, and programming (icalp 2016), 2016, pp. 8:1–8:13, doi:10.4230/LIPIcs.ICALP.2016.8.

[54] G. Grätzer, *General lattice theory*, Birkhäuser, Basel, 1978.

[55] R. Halin, *Über simpliziale Zerfällungen beliebiger (endlicher oder unendlicher) Graphen*, Mathematische Annalen **156** (1964), 216–225, doi:10.1007/BF01363288.

[56] ———, *Lattices of cuts in graphs*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **61** (1991), 217–230, doi:10.1007/BF02950767.

[57] ———, *Tree-partitions of infinite graphs*, Discrete Mathematics **97** (1991), no. 1-3, 203–217, doi:10.1016/0012-365X(91)90436-6.

[58] V. Hofmaier, *Kantenblöcke und Strukturbäume*, Bachelor's thesis, Universität Hamburg, 2016.

[59] J. Kneip, *Ends as tangles*, available at arXiv:1909.12628.

[60] ———, *Tangles and where to find them*, Doctorate thesis, Staats-und Universitätsbibliothek Hamburg Carl von Ossietzky, 2020. available at `https://ediss.sub.uni-hamburg.de/handle/ediss/8735`.

[61] F. Lazebnik and V. A. Ustimenko, *Explicit construction of graphs with an arbitrary large girth and of large size*, Discrete Applied Mathematics **60** (1995), no. 1-3, 275 –284, doi:10.1016/0166-218X(94)00058-L.

[62] H.-G. Leimer, *Optimal decomposition by clique separators*, Discrete Mathematics **113** (1993), no. 1-3, 99–123, doi:10.1016/0012-365X(93)90510-Z.

[63] C.-H. Liu, *Packing and covering immersions in 4-edge-connected graphs* (2015), available at arXiv:1505.00867.

[64] H. M. MacNeille, *Partially ordered sets*, Transactions of the American Mathematical Society **42** (1937), no. 3, 416–460, doi:10.2307/1989739.

[65] W. Mader, *Über n-fach zusammenhängende Eckenmengen in Graphen*, Journal of Combinatorial Theory, Series B **25** (1978), no. 1, 74–93, doi:10.1016/S0095-8956(78)80012-4.

[66] monkeymaths, *Two disjoint trees*, 2018, MathOverflow. `https://mathoverflow.net/q/308423` (version: 2018-08-16).

[67] N. Robertson and P. D. Seymour, *Graph minors. V. Excluding a planar graph*, Journal of Combinatorial Theory, Series B **41** (1986), 92–114, doi:10.1016/0095-8956(86)90030-4.

[68] ———, *Graph minors. X. Obstructions to tree-decomposition*, Journal of Combinatorial Theory, Series B **52** (1991), 153–190, doi:10.1016/0095-8956(91)90061-N.

[69] D. Seese, *Tree-partite graphs and the complexity of algorithms*, Fundamentals of computation theory, 1985, pp. 412–421, doi:10.1007/BFb0028825.

[70] P. D. Seymour, *Disjoint paths in graphs*, Discrete Mathematics **29** (1980), no. 3, 293–309, doi:10.1016/0012-365X(80)90158-2.

[71] R. E. Tarjan, *Decomposition by clique separators*, Discrete Mathematics **55** (1985), no. 2, 221–232, doi:10.1016/0012-365X(85)90051-2.

[72] C. Thomassen, *2-linked graphs*, European Journal of Combinatorics **1** (1980), no. 4, 371–378, doi:10.1016/S0195-6698(80)80039-4.

[73] A. W. Tucker, *Dual systems of homogeneous linear relations*, posted on 1957, 1–18, doi:10.1515/9781400881987-002.

[74] W. Tutte, *Graph theory*, Addison-Wesley, 1984.

[75] K. Wagner, *Über eine Eigenschaft der ebenen Komplexe*, Mathematische Annalen **114** (1937), 570–590, doi:10.1007/BF01594196.

[76] D. Weißauer, *Personal communication*, Hamburg, 2018.

[77] P. Wollan, *The structure of graphs not admitting a fixed immersion*, Journal of Combinatorial Theory, Series B **110** (2015), 47–66, doi:10.1016/j.jctb.2014.07.003.

Appendix A

Summary

This dissertation is concerned with the theory of tangles in abstract separation systems, which is part of the general field of structural graph theory. Tangles are a novel, universal tool to describe clusters in various structures, such as graphs, matroids or arbitrary data sets. They do so in an indirect way, built on the concept of separations of the given structure. They describe a cluster by deciding, for every possible way to 'sensibly' divide the structure into two parts, on which of the two sides the cluster is suspected to be located. I.e., they orient the separation given by the two parts towards the side containing the cluster. Certain consistency axioms on these orientations of separations ensure that the tangle indeed needs to point towards a cohesive region of the structure at hand.

The broad application of tangles to a variety of contexts was made possible by the general framework of abstract separation systems, in which tangles can be formulated in an abstract way. As such, results about tangles in this abstract framework can then directly be applied in the various contexts mentioned above. Consequently, also most parts of this thesis are formulated in the context of these abstract separation systems.

The results of this thesis are spread across three chapters. First, in Chapter 3, we are dealing with the task of finding concrete witnesses for the existence of a cluster described by a tangle. We provide a partial solution to a question by Diestel by showing that, given a tangle of a graph, we can always find a weight function on the vertices of the graph so that the given tangle chooses, for every separation of the graph, the side containing a higher total weight. We also analyse variations and generalizations of this question and consider the general question of how to witness the existence of a tangle via another structure in contexts other than graphs.

Moreover, we also provide a quantitative characterization of another type of highly connected structures in graphs, called agile sets, via specific minors.

In Chapter 4 we present results concerning the question of how to split up a structure according to their tangles. One of the two types of classic results about tangles, the *tree-of-tangles theorem*, states that, under certain conditions, a structure, such as a graph, can be split up in a tree-like way according to its tangles. We improve the two previously most general forms of this tree-of-tangles theorem significantly in multiple ways.

For example, we identify the key property needed for the various tree-of-tangle theorems to hold and thereby provide two simple, elementary lemmas which

allow to obtain the most relevant previously known tree-of-tangles theorems as corollaries. Moreover, the requirements of these lemmas are easy to check, and they allow us to obtain tree-of-tangles theorems for contexts where such theorems were previously not known. Additionally, we show that these lemmas even allow applications in contexts not covered by the framework of abstract separation systems, such as directed graphs, by also deducing a state-of-the-art tree-of-tangles theorem for tangles in directed graphs from one of the two lemmas.

We also show that trees of tangles can be constructed canonical, that is invariant under isomorphisms, in a more general context than previously known. While previously, a canonical tree-of-tangles theorem was only proven in the context of a *submodular universe of separations*, we prove such a theorem for *submodular separation systems*, a setup in which previously only the existence of a non-canonical tree-of-tangles theorem was known.

We additionally show the existence of tree-of-tangles theorems in infinite contexts, by providing, in the infinite context, a lemma analogue to the two simple lemmas mentioned above. This lemma can, as a side effect, be applied to the *separators* of an infinite graph instead of the separations. We show that the application to these separators gives us a canonical tree-like structure dividing up the graph according to its tangles which is of particular interest since both, a previously known non-canonical tree-of-tangles theorem for infinite graphs, and a previously known canonical theorem, which provides a structure called a *tree of tree-decompositions*, can be deduced from our theorem.

Moreover, our lemma about trees of tangles in infinite structures can also be used to obtain a tree-of-tangles theorem for the *edge-blocks* of a graph: the maximal subgraphs of an infinite graph which can not be separated by the deletion of some fixed number of edges.

We close Chapter 4 with showing that the other type of classic result from tangle theory, the *tangle-tree-duality* theorem, can in fact be used to prove a tree-of-tangles theorem. Thus, these two different classic types of theorems from tangle theory are not as independent as previously thought, since one of the two types of results can be used to obtain a result of the other type.

Finally, in Chapter 5, we investigate abstract separation systems as an object of its own interest. We not only provide examples of such systems which show that some of the previously defined notions on these systems, namely the notions of a submodular universe of separations and of a submodular separation systems mentioned above, indeed describe different objects, we also analyse their difference. This on the one hand leads to a decomposition theorem for submodular separation systems inside distributive universes, on the other hand we end up with a general question about certain families of finite sets, which we call *woven*. Our *unravelling problem*, which we state at the end of Chapter 5, is a simple to state problem about these woven sets. We analyse this problem, solve it in certain cases and provide solutions for some variations of that problem.

Appendix B

Deutsche Zusammenfassung

Diese Dissertation beschäftigt sich mit der Theorie von Tangles in abstrakten Teilungssystemen, einem Teilgebiet der strukturellen Graphentheorie. Tangles sind ein neues, universelles Werkzeug, um Cluster in verschiedenen Strukturen wie Graphen, Matroiden oder beliebigen Datenmengen zu beschreiben. Sie ermöglichen es, ein Cluster indirekt mit Hilfe des Konzepts der Teilungen einer Struktur zu beschreiben, indem ein Tangle für jede mögliche Art und Weise, die Struktur „sinnvoll" in zwei Teile zu zerteilen, entscheidet, in welchem der beiden Teile das Cluster vermutet wird. Mit anderen Worten, das Tangle „orientiert" die Teilung, die durch die zwei Teile gegeben ist, in Richtung der Seite, die das Cluster enthält. Bestimmte Konsistenz-Bedingungen an diese Orientierungen von Teilungen stellen sicher, dass das Tangle auch wirklich auf eine hoch-zusammenhängende Region der Struktur zeigt.

Das Konzept von Tangles kann abstrakt im Framework der abstrakten Teilungssysteme formuliert werden. Dies ermöglicht eine vereinheitlichte Anwendung von Tangles in einer Reihe von verschiedenen Bereichen, wie zum Beispiel den oben erwähnten. Dadurch können Erkenntnisse über Tangles in diesem abstrakten Framework direkt auf die verschiedenen Kontexte angewendet werden. Aus diesem Grund sind auch die meisten Teile dieser Arbeit im Kontext dieser abstrakten Teilungssysteme formuliert.

Die Ergebnisse dieser Arbeit werden in drei Kapiteln dargestellt. Zunächst beschäftigen wir uns in Kapitel 3 mit der Suche nach konkreten Strukturen, die die Existenz eines indirekt durch Tangles beschriebene Clusters garantieren. So präsentieren wir eine Teillösung zu einer Frage von Diestel, indem wir zeigen, dass man für ein gegebenes Tangle in einem Graphen immer eine Gewichtsfunktion auf den Ecken des Graphens finden kann, sodass das gegebene Tangle für jede Teilung des Graphens die Seite auswählt, die in Summe ein höheres Gewicht enthält. Zudem beschäftigen wir uns mit Variationen und Verallgemeinerungen dieser Frage, nämlich ganz generell damit, wie die Existenz eines Tangles auch außerhalb des spezifischen Kontextes von Graphen durch andere Strukturen bezeugt werden kann.

Außerdem charakterisieren wir über spezifische Minoren quantitativ die Existenz eines anderen Typen von hoch zusammenhängenden Strukturen in Graphen, sogenannten agilen Mengen.

In Kapitel 4 präsentieren wir mögliche Antworten auf die Frage, wie man eine gegeben Struktur entlang ihrer Tangles aufsplitten kann. Einer der zwei klassischen Sätze über Tangles ist der sogenannte „Tree-of-Tangles-Satz“, der besagt, dass man unter gewissen Bedingungen eine Struktur wie zum Beispiel einen Graphen baumartig entsprechend der Tangles zerlegen kann. Wir beweisen stärkere und allgemeinere Versionen der beiden bisher allgemeinsten Formen dieser Tree-of-Tangles-Sätze.

So identifizieren wir zum Beispiel eine zentrale Eingenschaft, die für die Gültigkeit solcher Tree-of-Tangles-Sätze notwendig ist, und verwenden diese, um zwei einfache, elementare Lemmata zu beweisen, die es ermöglichen, die relevantesten bisher bekannten Tree-of-Tangles-Sätze als Korollare zu folgern. Die Voraussetzungen dieser Lemmata sind zudem leicht nachprüfbar und erlauben es, Tree-of-Tangles-Sätze für bestimmte Typen von abstrakten Teilungssystemen zu beweisen, für die solche Sätze bisher nicht bekannt waren. Sogar eine Anwendung außerhalb des Frameworks der abstrakten Teilungssysseme ist möglich: Wir folgern einen Tree-of-Tangles-Satz für Tangles von gerichteten Graphen aus einem der beiden Lemmata.

Zudem zeigen wir, dass die Konstruktion eines Tree-of-Tangles in einer größeren Allgemeinheit als bisher bekannt kanonisch, also invariant unter Isomorphismen durchgeführt werden kann. Bisher war so eine kanonische Konstruktion nur für *submodulare Universen von Teilungen* bekannt, wir zeigen, dass eine solche kanonische Konstruktion auch in *strukturell submodularen Teilungssystemen* möglich ist. In diesen Systemen war bisher nur ein nicht-kanonischer Tree-of-Tangles-Satz bekannt.

Zudem entwickeln wir eine Variante der oben angesprochenen Lemmata für unendliche Strukturen, welche es uns ermöglicht, Tree-of-Tangles-Sätze auch für solche unendlichen Strukturen zu beweisen. Dieses Lemma können wir zudem auf die *Trenner* eines unendlichen Graphen anstelle seiner Teilungen anwenden, was es uns ermöglicht zu zeigen, dass man diese Trenner verwenden kann, um einen unendlichen Graphen kanonisch, baumartig, entsprechend seiner Tangles zu zerlegen. Diese Zerlegung ist insbesondere deswegen interessant, weil wir sie sowohl verwenden können, um einen bekannten, nicht kanonischen Tree-of-Tangles-Satz zu beweisen, als auch, um die Existenz eines kanonischen „Baums von Baumzerlegungen“ – ein anderes bekanntes Resultat über die Zerlegung unendlicher Graphen – zu folgern.

Zudem können wir unser Lemma über Trees-of-Tangles in unendlichen Strukturen verwenden, um einen unendlichen Graphen entsprechend seiner *Kantenblöcke* zu zerlegen. Ein Kantenblock ist hierbei ein maximaler Teilgraph, der nicht durch eine fixe Anzahl von Kanten getrennt werden kann.

Wir beenden Kapitel 4, indem wir zeigen, dass das andere klassische Resultat der Tangletheorie, der sogenannte *Tangle-Tree-Dualitätssatz*, dazu verwendet werden kann, einen Tree-of-Tangles-Satz zu beweisen. Dies zeigt, dass die bisher als unabhängig geltenden zwei zentralen Säulen der Tangletheorie – die Tree-of-Tangles-Sätze auf der einen und die Tangle-Tree-Dualitätssätze auf der anderen Seite – nicht so unabhängig voneinander sind wie man bisher dachte – schließlich kann man ein Resultat des einen Typs verwenden, um ein Resultat des anderen Typs zu folgern.

Im letzten Kapitel, Kapitel 5, dieser Dissertation untersuchen wir abstrakte Teilungssysteme als eigenständige Struktur. Zum einen zeigen wir, dass einige der existierenden verschiedenen Definitionen im Kontext dieser Systeme, nämlich zum Beispiel die oben erwähnten submodularen Universen von Teilungen und submodularen Teilungssysteme, in der Tat unterschiedliche Objekte beschreiben, indem wir entsprechende Beispiele konstruieren. Nachdem wir festgestellt haben, dass diese verschiedenen Definitionen in der Tat verschiedene Objekte beschreiben, untersuchen wir im nächsten Schritt, wie groß der Unterschied tatsächlich ist: verhalten sich die strukturell submodularen Teilungssysteme in gewissem Sinne „ähnlich" wie die submodularen Universen? Diese Analyse führt zum einen zu einem Zerlegungssatz für strukturell submodulare Teilungssysteme in distributiven Universen, zum anderen zu einer allgemeinen Frage über bestimmte Familien endlicher Mengen, die wir *verwoben* nennen. Unser *Entwirrungsproblem*, mit dem wir uns am Ende von Kapitel 5 beschäftigen, ist eine einfach zu formulierende Frage über diese verwobenen Mengen. Wir analysieren dieses Problem und lösen einige Varianten und Spezialfälle.

Appendix C

Publications related to this thesis

The following (pre-)publications are related to this dissertation:

Chapter 3

Section 3.1 is based on [37]. Section 3.2 is based on parts of [25]. Section 3.3 is based on parts of [24].

Chapter 4

Section 4.1 is based on [39]. Section 4.3 is based on [36]. Section 4.4 is based on parts of [42]. Section 4.5 is based on [43]. Section 4.6 is based on parts of [41].

Chapter 5

Section 5.1 is based on parts of [38]. Section 5.2 is based on [40].

Appendix D

Declaration of my contributions

The research in this thesis is based on work I performed with various co-authors, as can also be seen in Appendix C. In general, we performed the work in close collaboration and thus share an equal amount of work on both the research and the writing. I will now explain this in more detail, explaining which sections were developed with which co-authors, emphasizing where we do not share an equal amount of work as well as some highlights developed by me.

The results from Chapter 2 which are not taken from the existing literature are my own work.

Chapter 3

Section 3.1 covers joint work with Jakob Kneip and Maximilian Teegen. After having tried to find a decider for a tangle for over a year, Teegen came up with the idea to formulate the question in terms of linear programming. This led us three, in close cooperation, to an iterative and much more involved proof of Theorem 1 using Farkas' Lemma. This proof was simplified a lot by my finding of Tuckers Theorem in the literature and noticing that we can apply it. I also was the first to construct a version of the example in Example 3.1.7 and Theorem 3.1.8. The first two chapters of [37], on which the first two subsections of Section 3.1 are based, were drafted by us three together in close cooperation and then finalized by Jakob Kneip. The last subsection of Section 3.1 was drafted mostly by me, and then finalized mostly by Maximilian Teegen.

The paper [25] on parts of which Section 3.2 is based is joint work with Reinhard Diestel and Raphael Jacobs. Diestel set up the general questions answered in that paper and also came up with the notion of the resilience of a tangle and the question of whether tangles with higher resilience are more likely to have a decider. The results of [25] included in Section 3.2 have then been developed by me. Jacobs not only performed most of the writing of those results, but also developed additional characterizations of tangles with a decider in terms of local deciders and a duality of separations systems, which the interested reader can find in [25].

Section 3.3 is based on [24], which is joint work with Reinhard Diestel, Joshua

Erde, and Maximilian Teegen. Of the results of [24] presented in Section 3.3, I share no contribution on the original idea of shifting tangles and the other results presented prior to Theorem 6. However, both the proof of Theorem 6 as well as the idea of considering the underlying edge-tangles presented in Section 3.3.2 were developed just by Teegen and myself in close cooperation.

Section 3.4 is based on joint work with Jakob Kneip and Maximilian Teegen. We three started to look at Weißauer's Question 3.4.1 and developed the first results in Section 3.4, already including the generalizations. However, we were not able to come up with a quantitative characterization of graphs containing large agile sets, until I, revisiting the problem, found the paper [32] which allowed Maximilian Teegen and myself to prove Theorem 7. The first set of notes on which Section 3.4 is based was then also written together, however the final version as presented here is my own work.

Chapter 4

Section 4.1 is joint work again with Jakob Kneip and Maximilian Teegen. Teegen and myself started to work together on the question of whether there exists a theorem unifying both Theorem 1.1.3 and Theorem 1.1.2. We then developed a first version of Lemma 10, still formulated in terms of profiles, and with a more complex proof. That theorem already allowed us to establish Theorem 9. However, after presenting that proof to Kneip, he observed that both the statement and the proof of that theorem can be shortened to obtain Lemma 10 as presented in this thesis. The canonical version Lemma 11 was developed in close cooperation just by Teegen and myself. We two also drafted most of [39] on which Section 4.1 is based, whereas Kneip then wrote most of the final version.

The ideas in Section 4.2 were developed in joint work with Maximilian Teegen, with us two sharing roughly an equal amount of work. In particular, I came up with the idea to consider the relation of being weakly $\mathcal{P}$-nested.

Based on some rough notes of that work, I wrote Section 4.2 myself.

Lemma 12 was developed by Jakob Kneip, Maximilian Teegen and myself during the writing of [39].

Section 4.3 is joint work with Jakob Kneip. We share an equal amount of work on both the research and the writing. For example, while he came up with the question of whether a canonical version of Theorem 1.1.3 exists, I came up with the main ideas leading to our proof of Theorem 4.3.1. The paper [36] on which Section 4.3 is based, is in turn based on a corresponding section Jakob Kneip wrote for his dissertation.

The paper [42] on which Section 4.4 is based is joint work with Jakob Kneip and Maximilian Teegen. The results from [42] presented in Section 4.4 of this thesis were developed by Teegen and myself in close cooperation. In particular, I came up with both the definition of 'thinly splinters' and with the idea to consider the separators of the separations, which led to Theorem 16. Also, most of the drafting was performed by me, whereas the final version was written mostly by Teegen and myself in close collaboration. Kneip proved a version of our splinter lemma for *profinite separation systems*, which the interested reader can find in [42].

Section 4.5 is joint work with Jan Kurkofka and Maximilian Teegen. Kurkofka asked Teegen and myself, whether we know anything about a tree-of-tangles theorem for edge-blocks. In response to this, we applied Lemma 15 to develop Theorem 17, and also wrote a first draft of the proof. The relation to tree-cut decompositions was then observed by Kurkofka, who also wrote most of the final version of [43] on which Section 4.5 is based. Finally, we three together showed that Theorem 17 is equivalent to one important result from [17] by proving Theorem 18 in close cooperation.

Section 4.6 is again joint work with Jakob Kneip and Maximilian Teegen. The original question was set by Nathan Bowler and Joshua Erde, who asked whether it is possible to obtain a tree of tangles from the general tangle-tree duality Theorem 4.6.1. In response to this, Teegen and myself developed a first such proof, which runs along the lines of the proof without using Lemma 4.6.10 presented in Section 4.6.3. Kneip then observed that Lemma 4.6.10 can not only be used to simplify that proof, but also to obtain a condition on the degrees of the resulting tangles, as shown in Section 4.6.4. He also wrote most of the sections of [41] on which Sections 4.6.3 and 4.6.4 are based. The generalization to trees of tangles distinguishing the profiles efficiently (Sections 4.6.5 and 4.6.6), as well as to tangles of different orders (Section 4.6.7), was developed just by Teegen and myself, including the writing of the corresponding sections of the paper [41], on which the corresponding sections of this thesis are based.

Chapter 5

Section 5.1 is based on [38], which is joint work with Jakob Kneip and Maximilian Teegen. However, the parts from that paper presented in this thesis are joint work with Maximilian Teegen only, which we performed in close cooperation. In particular, I came up with the idea of dependency digraphs, as well as the proof of Theorem 5.1.7. We also wrote the corresponding sections of the paper [38], on which Section 5.1 is based, together. The paper [38] includes an additional section developed just by Kneip and Teegen about a version of the Dedekind-MacNeille completion for separation systems which are submodular in the sense of Section 5.2.6, which is not part of this thesis.

[40] on which Section 5.2 is based is also joint work with Jakob Kneip and Maximilian Teegen. Our research started with Teegen and myself showing Lemma 5.2.8. This then resulted in Kneip developing the other results in Section 5.2.4 by showing that we may always assume our order-function to have distinct values. He also wrote the corresponding section of [40] on which Section 5.2.4 is based. The research that led to Sections 5.2.5 and 5.2.6 was performed just by Teegen and myself, and the corresponding sections of [40], on which these sections are based, were also written by us two together. The final construction of Section 5.2.5 is mine, including the idea of considering a regular graph with high girth.

Appendix E

Acknowledgement

First, I would like to thank my supervisor Reinhard Diestel. You not only always supported my research with ideas and advice, you also formed my fascination for graph theory during a bachelors course. Thanks a lot for the opportunity to do a PhD under your supervision.

I would also like to thank my (former) office-mates and co-authors Jakob Kneip and Maximilian Teegen. We worked together a lot during our doctoral studies. Thank you not only for the great working experience, but also for great discussions not only about maths but also about anything else, it was always a lot of fun working with you together on a mathematical problem, and during this four years you really became friends for me.

Moreover, I would like to thank my other co-authors, Josh Erde, Jan Kurkofka, and Raphael Jacobs; it was a lot of fun working with you together on the corresponding projects.

I would also like to thank the whole discrete mathematics research group in Hamburg, for creating such a great working climate, not only during office hours, but also way beyond. Our game nights were always a lot of fun!

And finally, I would like to thank my family for all the love and support during the four years of my PhD.

Appendix F

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.

Christian Elbracht