Tangles and where to find them

Dissertation

zur Erlangung der Würde des
Doktors der Naturwissenschaften

der Fakultät für Mathematik, Informatik
und Naturwissenschaften,

Fachbereich Mathematik der

Universität Hamburg

vorgelegt von

Jakob Kneip

Hamburg, 2020

Vorsitzender der Prüfungskommission:
Prof. Jens Struckmeier (Universität Hamburg)

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

Zweitgutachter:
Prof. Paul Wollan (University of Roma "La Sapienza")

Datum der Disputation:
22. Oktober 2020

Contents

Chapter 1

Introduction

1.1 Tangles: A History

Connectivity is among the most central and classical topics in graph theory: there are myriads of notions and results regarding when exactly a set of vertices in a graph, or the graph itself, is 'well-connected' in some sense, and the properties exhibited by such well-connected graphs or parts thereof have been deeply analysed for many of these possible ways of defining connectivity.

Among the most classical, and most widely applied, results concerning connectivity is Menger's Theorem[1], which asserts that the minimum order of a separation separating two given vertex sets is equal to the maximum number of disjoint paths connecting these two sets. In this way Menger's Theorem creates a link between a *separating* statement and a *connecting* statement: if the graph can be separated in a certain way, then there is no large set of paths or connections; and conversely, if the graph cannot easily be separated, then we always find a family of paths of adequate size.

The highly connected structure in Menger's Theorem takes the form of two arbitrary vertex sets joined by many disjoint paths. There are various other notions of cohesive objects or regions: *cliques*, which are complete subgraphs, and more generally clique minors; *blocks*, which are inclusion-maximal two-connected subgraphs, and their generalisation of k-blocks (cf. [5]); complete bipartite subgraphs; and grids or grid minors, to name just a few. All of these have in common that in some way the vertices belonging to one such structure are difficult to separate from each other in the host graphs – although what exactly this should mean depends on the structure at hand.

Blocks in particular illustrate another common concept concerning connectivity: it is an easy exercise to show that the blocks of a graph G induce a tree structure on G. This property can be found in many manifestations of connected structures: if G contains many highly connected objects in different places, then typically these can be separated from each other in a tree-like fashion. This is true for the more general concept of k-blocks as well: as was shown in [5], any graph admits for each k a tree decomposition which separates its k-blocks.

The general assertion of Menger's Theorem that the absence of separability

[1] This classical result, as well as definitions of all concepts used but not explicitly defined in this section, can be found in [10].

implies the presence of connectivity can be observed in the aforementioned notions of highly connected structures as well. For instance, a classical result by Robertson and Seymour [40] asserts that if a graph G cannot efficiently be cut into a tree-like structure, i.e. if it has no tree decomposition with small parts, then G must contain a large grid minor. In contrast to Menger's Theorem this result does not achieve equality of the parameters though: to force the existence of a $k \times k$-grid minor in a graph it is necessary to forbid the existence of a tree decomposition into parts of size $f(k)$ for some function f whose optimal values are still subject of ongoing research. (See, for instance, [7].)

In their Graph Minor Project [41], Robertson and Seymour proposed the new idea of *tangles* as a tool to unify all of the above-mentioned notions of connectivity into one, and treat these in a single framework. The core idea of tangles is that although the exact definitions of various highly cohesive objects in graphs may vary, they all share the property that it is not possible to cut right through them with a separation of low order. In other words: for each low-order separation of a graph G, any given well-connected structure ought to lie mostly one one side of that separation.

The observation of Robertson and Seymour then was that to study these well-connected structures it often suffices not to know that structure explicitly, but only the orientation of the low-order separations it defines. These orientations are the *tangles* of a graph, and they can be analysed independently of any particular concrete notion of connectivity.

Formally a *separation* of a graph $G = (V, E)$ is a pair (A, B) of subsets of V with $A \cup B = V$ and such that G has no edge incident with both a vertex in $A \setminus B$ and one in $B \setminus A$. For an integer k a *k-tangle* of G consists of exactly one of (A, B) and (B, A) for each separation of G with $|A \cap B| < k$, and has the property that there are no $(A_1, B_1), (A_2, B_2), (A_3, B_3)$ in the tangle whose left sides cover G in the sense that $G = G[A_1] \cup G[A_2] \cup G[A_3]$.

It is not hard to see that each of the well-connected structures above, if sufficiently large, defines a k-tangle in G by letting the tangle contain out of (A, B) and (B, A) that pair for which the well-connected structure (or the majority thereof) is contained in the right-hand side.

These tangles still enable one to prove both types of typical connectivity result: if a graph is not well-connected in the sense that it does not have a k-tangle, then it can efficiently be cut into a certain tree structure. Moreover, as with blocks, the tangles of a graph can be separated from each other in a tree-like fashion as well. We call the first type of result a *tangle-tree duality* result, and the second one a *tree-of-tangles* result; both were established in the inaugural tangle work [41].

This approach via tangles, and the shift of paradigm coming with it, have one more upside: they bring more algebraic tools to the connectivity table. Tangles are made up of separations of a graph, and these separations exhibit a rigid structure: by defining a partial order through $(A, B) \leqslant (C, D)$ whenever $A \subseteq C$ and $B \supseteq D$, the set $\vec{S}(G)$ of all separations of a graph G becomes a lattice. Moreover, the function $(A, B) \mapsto |A \cap B|$, which is integral to the definition of a k-tangle, is submodular on this lattice. Submodular functions are a regular guest in connectivity theory: they come up when dealing with cuts of graphs, but also in matroids. (See [37] for a variety of examples of submodular functions in graph theory and related fields.) Structures defined by submodular functions have not only some desirable algorithmic properties – which are of no concern to us in this work – but also facilitate some standard optimization arguments in combinatorics, which make tangles especially pleasant to work with.

1.2 Tangles: Modern Theory

A subtle difference between tangles and their predecessor connectivity notions is that tangles do not describe a cohesive structure by listing the vertices it consists of, as a clique or a grid minor would: instead, the tangle sets up a system of 'pointers' in the form of oriented separations which tells one *where* the dense structure is, but makes no membership assertion for any single vertex of the graph. This shift of paradigm away from individual vertices, together with the observation that separations of graphs form lattices, gave rise to the idea that the vertices of a graph, or indeed the graph itself, might not be needed to formulate the theory of tangles.

Indeed, it turned out that all that is needed from the separations of a graph to formulate and study tangles is the partial order on its separations, together with the map $(A, B) \mapsto (B, A)$ which maps each separation to its 'inverse'. Thus *abstract separation systems* were born: partially ordered sets $\vec{S} = (\vec{S}, \leqslant)$ together with an involution $^*\colon \vec{S} \to \vec{S}$ which is order-reversing – which is to say that if two elements of $\vec{S}$ are comparable, then their images under * shall be comparable in the opposite direction. This involution can be expressed efficiently in notation by writing $\overleftarrow{s}$ for the inverse $(\vec{s})^*$ of $\vec{s} \in \vec{S}$; using this notation we then have $\vec{r} \leqslant \vec{s}$ if and only if $\overleftarrow{r} \geqslant \overleftarrow{s}$. If such an abstract separation system is a lattice we call it a *universe (of separations)*. The separations of a graph form such a universe of abstract separations. A *tangle* in this abstract set-up then is a set consisting of exactly one of $\vec{s}$ and $\overleftarrow{s}$ for each $\{\vec{s}, \overleftarrow{s}\} \subseteq \vec{S}$ and satisfying certain consistency requirements. These requirements can be handled quite flexibly: by forbidding the tangles of $\vec{S}$ to include certain configurations of separations one can not only emulate graph-theoretical tangles in this setting but also define and study new types of tangles by varying which configurations one forbids. Those tangles of graphs that come from well-connected objects are easily seen to obey certain consistency axioms: no two comparable separations $(A, B) \leqslant (C, D)$ oriented by the dense object in the graph can be oriented 'inconsistently', that is, pointing away from each other, since clearly the well-connected structure cannot lie in both A and D.

With the shift to abstract separations and tangles, and the theoretical upsides of this new approach, the floodgates opened and advances in tangle theory were made in quick succession ([12, 13, 16, 17, 19, 20]). First and foremost, forgetting that the separations at hand originate from a graph allows one to formulate a theory of connectivity that is applicable not only to graphs, but to many other combinatorial structures as well, such as matroids or bipartitions of sets. The most far-reaching consequence of this is that tangles as a tool become available to entirely different fields of science: for instance, [23] proposed some ways in which tangles can be used in clustering algorithms in computer science, both for the traditional task of finding clusters in graphs as well as for identifying typical 'mindsets' displayed by the participants of some questionnaire. Another upside of the abstract view on tangles is that this algebraic setting allows for many clean and elegant proofs of classical results, whose original proofs sometimes get bogged down with details that turn out to be unnecessary.

For both of the two archetypical tangle results, the tree-of-tangles theorems and tangle-tree duality, the abstract setting offers the strongest, cleanest, and most widely applicable versions. The tree-of-tangles theorem established in [17] finds a tree structure which efficiently distinguishes all the given tangles of an abstract separation system, which includes the graph-theoretical tangles from [41] as a special case. However, [17]'s new tree-of-tangles theorem is not only more general than [41]'s, but also brings an entirely new quality to the table: the tree structure it finds is *canonical*, which means that it can be found using only invariants of the underlying structure. This is in stark contrast to the approach of [41], whose output tree structure can vary wildly depending on choices made arbitrarily during its construction.

The improvements regarding tangle-tree duality are even more impressive compared to the graph-theoretical results: [19] established the *unified tangle-tree duality theorem* which for a very general definition of tangle proves that a given separation system admits either a tangle or a tree-structure which demonstrates that it cannot have a tangle. Utilising the freedom of choosing which type of tangle to plug into the theorem one can quickly and easily obtain tangle-tree duality results for tangles arising from various connectivity notions, each of which gives rise to a differently shaped tree structure as the dual object of its particular tangles. Moreover one can turn this process around: if one is able to describe a certain kind of tree structure through the configurations of separations it consists of, then the unified tangle-tree duality theorem delivers a custom tangle variety that is dual to that tree structure. Finally, one can even use the unified tangle-tree theorem to establish other results: by defining a type of tangles of which one knows that it cannot exist in a given separation system, one guarantees that the duality theorem outputs a tree structure of this system, the shape of which one can control via the type of tangles one considers. In this way one can prove theorems which, on the surface, have nothing to do with duality. (We shall see more of this in Section 4.3.)

Despite these theoretical successes, the modern theory of abstract tangles on the whole is still young and fresh and subject of many ongoing investigations. In this work we will carry out some of these investigations throughout every facet of tangle theory, and answer some open questions.

1.3 Tangles: New Contributions

The results presented in this work can also be found in [2, 24–27, 33], except for those which we will point out explicitly as unpublished both here and in their chapter's introductions. The results not published outside of this work at time of writing are in Sections 3.4 to 3.6, 4.1.4, 4.1.7, 4.1.8, 4.2.2, and 5.2.5.

We give an overview of the definitions and notational conventions of tangle theory in Chapter 2. Following that we begin our journey into tangle theory proper. For this we begin by studying the general structural properties of separation systems. The most important and most frequently encountered types of separations are separations of graphs and derivatives thereof. The most prominent alternative to separations of graphs are bipartitions of sets, which occur naturally as cuts of graphs. Much of the intuition about abstract separation systems stems from these classes of examples, and quite a few abstract definitions are modelled on them. A natural question is thus just how different an abstract separation system can be from one of these examples; or put in another way, which structural properties an abstract separation systems needs to display to be representable by separations of a graph or of bipartitions of a set.

In Chapter 3 we answer this question and characterise the separation systems coming from graphs or bipartitions by their combinatorial properties. For instance, our characterisation of those universes of separations that come from bipartitions of some ground-set reads as follows:

Theorem 4. *A universe of separations $\vec{U}$ can be strongly implemented by bipartitions of sets if and only if it is distributive and fastidious.*

Here a universe $\vec{U}$ of separations is *fastidious* if $\vec{r} \leqslant \overleftarrow{r}$ for some $\vec{r} \in \vec{U}$ implies $\vec{r} \leqslant \vec{s}$ for all $\vec{s} \in \vec{U}$. Using the characterisations from Sections 3.2 and 3.3 one can tell whether the separation system one is dealing with is effectively one of those prominent types.

In the remainder of Chapter 3 we answer a variety of other general questions about the properties of separation systems; the results obtained there are all unpublished. Given a separation system $\vec{S}$, one can define a graph G_S using the separations of $\vec{S}$ as vertices, and joining two separations with an edge whenever they have no comparable orientations; we call this the *crossing graph* of $\vec{S}$. (Separations with comparable orientations are said to be *nested*, and those without to *cross*.) In Section 3.4 we find out just how much information can be gained about $\vec{S}$ from the shape of its crossing graph G_S, and use our insights to give a more efficient version of one of Section 3.2's representation theorems concerning bipartitions of sets, a close cousin of Theorem 4. We then analyse the class of graphs that can occur as a crossing graph G_S. These crossing graphs turn out to be fairly common – every graph is the crossing graph of a suitable separation system:

Theorem 8. *For every graph G, not necessarily finite, there is a separation system S such that G is isomorphic to G_S.*

We then turn our attention to submodularity. Both of the archetypical examples of separation systems, graph separations and bipartitions in the form of cuts, come with a natural submodular order function. For a graph $G = (V, E)$, the functions mapping a separation (A, B) to $|A \cap B|$ or a cut (A, B) to the number of edges with endpoints in both A and B can be used as 'efficiency measure' or 'cost function' for those separations or cuts: the lower the value of those function on some (A, B), the more efficient that separation or cut of G is at disconnecting G. We call such maps *order functions*. Both of these functions are submodular on the respective lattice of separations. Consequently the separation system $\vec{S}_k$ consisting of all separations/cuts of G whose value of this function is below some threshold k has the structural property that for any two separations in $\vec{S}_k$ at least one of their pairwise join and meet also lies in $\vec{S}_k$. This latter property is a purely structural one and can therefore be defined without the need for such a submodular order function: a separation system $\vec{S}$ inside some lattice is (structurally) submodular if at least one of $\vec{r} \wedge \vec{s}$ or $\vec{r} \vee \vec{s}$ lies in $\vec{S}$ for all $\vec{r}, \vec{s} \in \vec{S}$. Structural submodularity is one of the central properties in tangle theory, and in fact makes both tree-of-tangles and tangle-tree duality results attainable without requiring additional structure ([16]). This submodularity property itself is therefore a worthy object of study.

In Section 3.5 we aim to rid ourselves of the technical annoyance that in order for some separation system $\vec{S}$ to be submodular we need a surrounding lattice structure in which we can express and verify this: for $\vec{S}$ to be submodular it is not in general necessary or typically the case that $\vec{S}$ itself is a lattice. We show that if $\vec{S}$ is submodular in the sense that any two elements of $\vec{S}$ have either a pairwise supremum or infimum in $\vec{S}$, a formulation which makes do without a surrounding lattice structure, then we can embed $\vec{S}$ into a suitable lattice in such a way that $\vec{S}$ is then submodular inside that lattice measured in the original sense:

Theorem 11. *For every separation system $\vec{S}$ there are a universe $\vec{U}$ and a map $f\colon \vec{S} \to \vec{U}$ that is an isomorphism of separation systems between $\vec{S}$ and its image in $\vec{U}$, with the property that $f(\vec{t}) = f(\vec{r}) \vee f(\vec{s})$ if and only if $\vec{t}$ is the supremum of $\vec{r}$ and $\vec{s}$ in $\vec{S}$, and likewise $f(\vec{u}) = f(\vec{r}) \wedge f(\vec{s})$ if and only if $\vec{u}$ is the infimum of $\vec{r}$ and $\vec{s}$ in $\vec{S}$.*

In particular if $\vec{S}$ is submodular as a poset then $f(\vec{S})$ is submodular as a separation system inside $\vec{U}$.

Concluding Chapter 3, in Section 3.6 we further investigate submodularity. The motivating question of that section is the following: if $\vec{S}$ is a structurally submodular separation system, is there always some $\vec{s} \in \vec{S}$ whose deletion from $\vec{S}$ leaves $\vec{S}$ submodular? A positive answer to this would facilitate some elegant proofs by induction, an example of which we give in a later chapter in Section 4.1.8. As a first step we show that the desired assertion is true for all $\vec{S}$ that take the form $\vec{S}_k$ for some submodular order function on the surrounding lattice, as defined above. For this we show that if $|\cdot|$ is a submodular order function on some lattice of separations, then $|\cdot|$ can be modified so as to be injective while maintaining submodularity and keeping pre-existing strict inequalities. The separation system $\vec{S} = \vec{S}_k$ can then be 'unravelled' by successively deleting the unique $\vec{s}$ which maximises $|\vec{s}|$ from $\vec{S}$: this keeps $\vec{S} = \vec{S}_k$ submodular throughout, until we have reduced it to the empty set.

To make a given submodular order function injective we first show that every finite universe of separations admits a submodular and injective order function. An appropriately scaled sum of the given order function and this 'tiebreaker' will then be injective and submodular while preserving strict inequalities.

Curiously Robertson and Seymour employed a similar 'tiebreaker' technique in [41] to modify the orders of separations of a graph so as to be distinct while essentially keeping submodularity. As their application of these tiebreakers they established the first ever tree-of-tangles theorem for tangles in graphs:

Theorem 14 ([41])**.** *Every graph has a tree-decomposition displaying its maximal tangles.*

This tree-of-tangles theorem by Robertson and Seymour is also the starting point of Chapter 4, in which we present our results in finite tangle theory. The original proof of the tree-of-tangles theorem given in [41] is, all told, about eight pages long, and our first mission in Chapter 4 will be to employ the tools and strategies of the modern abstract tangle theory to give a new proof that is as short and elegant as possible. We achieve this by establishing the following result:

Theorem 18 (Splinter theorem)**.** *Let U be a universe of separations and $\mathcal{A} = (A_i)_{i \leqslant n}$ a family of subsets of U. If $\mathcal{A}$ splinters then we can pick an element a_i from each A_i so that $\{a_1, \dots, a_n\}$ is nested.*

The proof of Theorem 18 takes less than half a page, and the condition imposed on the family $\mathcal{A}$ to splinter is both simple to define and easy to check. The verification that Theorem 18 does indeed imply Theorem 14 is carried out in Section 4.1.3 and is straightforward: as the set A_i we take those separations of the graph G that are of lowest order in the symmetric difference of the i-th pair of tangles. These sets are easily seen to satisfy the splinter condition, and a set of pairwise nested separations containing one element of each A_i gives rise to the desired tree-decomposition as in Theorem 14.

However, the purpose of Theorem 18 is not only to give a new shortest proof of Theorem 14. By following the same general strategy one can show that Theorem 18 implies virtually every other tree-of-tangles theorem, including some of the more recent results in tangle theory. Moreover Theorem 18 can be used to establish new tree-of-tangles theorems in a variety of settings. For instance, we establish a tree-of-tangles theorem for clique separations by utilising the flexibility in choosing the sets A_i and the simplicity of the splinter condition. Since the clique separations of a graph – those separations (A, B) for which $A \cap B$ is a clique – do not form a lattice structure, this tree-of-tangles theorem could not have been obtained using the traditional tools of tangle theory. This demonstrates that despite its short and elementary proof Theorem 18 is quite powerful.

We also prove a variation of Theorem 18 that is *canonical*, i.e. which uses only invariants of the given separation systems to find the desired nested set of separations. This variation, too, can be used both to recover known results with more compact proofs and to establish new canonical tree-of-tangles theorems.

In Section 4.1.4 we give an even more abstract version of Theorem 18, based on the observation that Theorem 18 and its proof only use *some* of the properties of separation systems: namely the information which separations are nested and which cross, and the interplay of these relations with taking joins or meets. This information can be expressed using only the reflexive and symmetric relation given by 'being nested', a concept complementary to the crossing graphs G_S. This further abstraction in principle makes Theorem 18 applicable outside of tangle theory, and also serves to make the proof shorter yet, coming out at just eight lines. This addendum to Theorem 18 is not published outside of this work.

In Section 4.1.7 we present another result not found outside of this work. The canonical variant of Theorem 18 is unable to prove a canonical tree-of-tangles theorem for those separation systems that are only structurally submodular, but not of the form $\vec{S}_k$ for some order function. (Theorem 18 does imply a non-canonical tree-of-tangles theorem for those separation systems though, see Theorem 16.) Using a different and more hands-on approach we show that these separation systems give rise to canonical trees-of-tangles after all:

Theorem 25. *Let $\vec{U}$ be a finite universe of separations, $\vec{S} \subseteq \vec{U}$ submodular, and $\mathcal{P}$ a set of profiles of S. Then there is a nested set $N = N(\mathcal{P}) \subseteq S$ which distinguishes $\mathcal{P}$. This $N(\mathcal{P})$ can be chosen canonically: if $\varphi\colon \vec{U} \to \vec{U'}$ is an isomorphism of universes, then $\varphi(N(\mathcal{P})) = N(\varphi(\mathcal{P}))$.*

The 'profiles' mentioned in Theorem 25 are a general-purpose type of tangle, whose only consistency conditions are precisely those that are needed to make a tree-of-tangles theorem possible. Virtually all sensible types of tangles are such profiles, and the tree-of-tangles theorems we present here are often formulated for profiles to make them applicable to the widest possible range of tangles.

Moving on from tree-of-tangles theorems, we turn our attention to the second archetypical result of tangle theory: the tangle-tree duality theorem. As previously mentioned, [19] established the *unified tangle-tree duality theorem*, which can be applied very flexibly and used to obtain tangle-tree duality theorems in a large variety of settings:

Theorem 1 (Tangle-tree duality theorem [19])**.** *Let U be a finite universe of separations, $S \subseteq U$ 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 $\mathcal{F}$-separable. Then precisely one of the following holds:*

- *there is an S-tree over $\mathcal{F}$;*
- *there is an $\mathcal{F}$-tangle of S.*

However, so far Theorem 1 is somewhat detached from the rest of tangle theory due to the fact that almost all definitions introduced and techniques employed in [19] are exclusive to tangle-tree duality and not applicable or usable elsewhere. Moreover the proof of Theorem 1 unfortunately requires a lot of technical set-up and preliminary results, and is itself somewhat technical. In Section 4.2 we seek to alleviate these concerns by giving two new proofs of Theorem 1. Both of these new proofs strive to be more enlightening as well: the original proof of Theorem 1 proceeds by a technical induction, whereas these new proofs are more constructive and hands-on. The upsides offered by the proof methods presented here are not just in building a better intuition, however. They also yield a significant weakening of one of the assumptions of Theorem 1 (namely, of S being $\mathcal{F}$-separable, which we shall not define here), thereby strengthening the theorem and making it more widely applicable. Indeed, one application of Theorem 1 found in [24] requires our strengthened version of Theorem 1 and would not be possible with the theorem as stated above. For this reason the first of our two proofs is also included in [24]; the second, presented here in Section 4.2.2, is exclusive to this work. This second proof, while not altogether shorter than Theorem 1's original one, aims to cut down on the amount of technical preliminary results needed to establish Theorem 1, and thereby make its proof more self-contained.

We demonstrate the flexibility of Theorem 1 in Section 4.3 by presenting an unexpected application thereof: we prove a tree-of-tangles theorem by using the tangle-tree duality theorem. Concretely we show the following:

Theorem 28. *Let $\vec{S}$ be a submodular separation system. Then S contains a nested set that distinguishes the set of regular profiles of S.*

This Theorem 28 is a slightly weaker and, more significantly, non-canonical version of Theorem 25 above. The result itself is therefore nothing special, but the method by which we prove it is: we carefully craft a set $\mathcal{F}$ for which we know that there is no $\mathcal{F}$-tangle of S, and for which we therefore must get the first outcome in Theorem 1, an S-tree over $\mathcal{F}$. By our choice of $\mathcal{F}$ this S-tree over $\mathcal{F}$ will then already represent the claimed nested set in Theorem 28, with no further modifications required. Remarkably, in this way we obtain a tree-of-tangles for one kind of tangle by constructing a tree structure witnessing the non-existence of another carefully chosen type of tangle. Thus evidently the two archetypical tangle theory results are more closely related than previously thought.

We conclude Chapter 4 by returning to the roots of tangle theory: k-tangles in graphs. As explained above, the intuition behind k-tangles in graphs is that they capture well-known highly connected structures by the orientations of low-order separations those structures define. The way in which a sufficiently large and well-connected object in a graph gives rise to a k-tangle is then that for any separation (A, B) with $|A \cap B| < k$, exactly one of A and B will contain the vast majority of that well-connected object. Out of (A, B) and (B, A) the tangle then contains (A, B) rather than (B, A) if and only if B is the side containing most of the highly cohesive object.

As a part of this intuition about tangles in graphs we think of B as the 'big side' of a separation (A, B) in a tangle, and of A as the 'small side'. So far it has been an open question whether this intuition can be made concrete: is it true that for every k-tangle in a graph there is some set X of vertices such that a separation (A, B) with $|A \cap B| < k$ lies in that tangle if and only if B contains a majority of the vertices of X? This natural question was raised by Diestel in [17], and an affirmative answer would substantiate the notion that tangles arise from cohesive objects by orienting separations towards the majority of those objects.

While we cannot quite give a positive answer to Diestel's question, we can prove a fractional version of it:

Theorem 30. *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 $U \subseteq V$.*

In other words, for every k-tangle of a graph there is a *weighted* set X of vertices, or equivalently a multiset X, such that B is the 'big side' of a separation (A, B) in this tangle if and only if B contains a strict majority of the (weighted) vertices in X. We show that this result extends to tangles of hypergraphs as well as to some other types of tangles in graphs. However, we also give an example of a type of tangles for which Theorem 30 fails, demonstrating that Theorem 30 cannot be made applicable to arbitrary tangles in (hyper-)graphs.

In our final chapter, Chapter 5, we investigate the tangle theory of infinite separation systems and its applications to infinite graphs. Extending results for finite separation systems to infinite ones is usually a difficult task: most tree-of-tangles and tangle-tree duality theorems as well as their proofs work by induction, or use the finiteness of the separation system in some other implicit way such as considering the maximal elements of some partial order. The proofs presented in Chapter 4 of this work are no exception to this: easy extensions to the infinite setting are available for none of the results outlined above.

At the same time the definition of tangles in graphs can be used verbatim in infinite graphs as well, and infinite graph theory is an active and deep field within graph theory. It therefore would be strongly desirable to lift some of the results concerning finite separation systems to infinite ones – if not to arbitrary separation systems then at least to ones with certain combinatorial properties such as the separation systems of infinite graphs.

The method by which results for finite graphs are usually lifted to infinite graphs is *compactness*. The basis of this principle is that an infinite graph G can be 'built up' from its finite subgraphs or finite minors. If one can solve the problem at hand for each of these finite subgraphs, and do so in a way that the solutions for different subgraphs are in some sense 'compatible', then one can hope that these finite solutions lift to a common solution of the infinite problem in G. Solutions for the finite subgraphs of G usually can be obtained by applying the already established finite version of the theorem one intends to prove for infinite graphs.

This compactness method for infinite graphs is facilitated by the fact that every infinite graph G is uniquely determined by the family of its finite subgraphs: if one knows $G[X]$ for every finite set X of vertices of G, then one knows all of G. With this fact and the correct technical definition of 'compatible solutions' one can then ensure that any family of compatible solutions for the finite subgraphs of G indeed lifts to a solution for G. Most importantly for us though, this observation is true for separations of G as well: if (A, B) is a separation of G then for every finite set X of vertices of G the restriction of (A, B) to X is a separation of $G[X]$. Moreover one can recover (A, B) from the family of these separations of the $G[X]$. Indeed, (A, B) is a separation of G if and only if each of its finite restrictions is a separation of the corresponding subgraph. Thus every family of separations of the finite subgraphs that is compatible in a certain technical sense indeed gives rise to, or comes from, a separation (A, B) of G. Therefore the compactness method is available for use in the tangle theory of infinite graphs as well.

The above observation shares some concepts with the notion of *profinite topological spaces* from general topology. (See [18] for a more in-depth discussion.) Indeed, one can define a topology on the separation system $\vec{S} = \vec{S}(G)$ in a straightforward manner: for each finite set X of vertices one equips $\vec{S}(G[X])$ with the discrete topology. Then $\vec{S}$ is a (closed) subspace of the product of all these $\vec{S}(G[X])$, where each separation $(A, B) \in \vec{S}$ of G is to be understood as an element of this product by way of the family of its restrictions to the $G[X]$. This topology on $\vec{S}$ has many useful properties: it is compact Hausdorff, and furthermore a set of separations (such as a tangle) is closed in $\vec{S}$ if and only if that set can be recovered exactly from the family of subsets of $\vec{S}(G[X])$ its restrictions induce. This topology on $\vec{S}$ therefore enables us to express which sets of separations and which tangles of G are amenable for the compactness method.

Our first object of study in Section 5.1 will be the ends of infinite graphs. Ends of graphs are not in general recognised as well-connected objects or dense regions. However, it is well-known that an end is infinitely dominated if and only if it 'lives in' a subgraph that is a subdivided infinite clique. The latter certainly is a well-connected object, and hence some ends can indeed be seen as a highly cohesive region.

Every end of an infinite graph, however, gives rise to a tangle in that graph regardless of its domination: for an infinite graph G each end of G defines a tangle of the set $\vec{S}_{\aleph_0}$ of all separations (A, B) of G with $A \cap B$ finite by containing such an (A, B) if and only if that end has a tail in B. We call the tangles arising in this way the *end tangles* of G. Our view on ends will be exclusively through the lens of their end tangles.

Some research into these end tangles has already been carried our in [11]. One of the main results of that work makes a connection between two of the notions discussed above: an end of an infinite graph G defines a tangle that is closed as in the subspace $\vec{S}_{\aleph_0}$ of the above topology if and only if that end is infinitely dominated, and hence constitutes a well-connected object in the classical sense.

Each end tangle of a graph G orients all separations (A, B) with $|A \cap B| < \infty$. Thus in particular for any integer k that end orients all separations with $|A \cap B| < k$ and hence defines a k-tangle of G. These k-tangles, too, may or may not be closed as subsets of $\vec{S}$, and one can attempt to draw a connection similar to the one above between this closedness for certain values of k and the combinatorial properties of the end in question. First of all it is true that if an end tangle is closed in $\vec{S}_{\aleph_0}$ then so is the k-tangle induced by that end for each k, for the simple topological reason that the set $\vec{S}_k$ of all separations with $|A \cap B| < k$ is always a closed subset of $\vec{S}$.

The connection we draw in Section 5.1.4 is the following:

Theorem 36. *Let τ the end tangle induced by an end ω of G. Then the following statements hold:*

(i) *τ is closed in $\vec{S}$ if and only if $\operatorname{dom}(\omega) = \infty$.*

(ii) *τ is not closed in $\vec{S}$ but $\tau \cap \vec{S}_k$ is closed in $\vec{S}_k$ for every $k \in \mathbb{N}$ if and only if $\deg(\omega) = \infty$ and $\operatorname{dom}(\omega) < \infty$.*

(iii) *$\tau \cap \vec{S}_k$ is not closed in $\vec{S}_k$ for some $k \in \mathbb{N}$ if and only if $\deg(\omega) + \operatorname{dom}(\omega) < \infty$.*

Here $\deg(\omega)$ and $\operatorname{dom}(\omega)$ are the vertex degree and number of dominating vertices of an end ω, respectively. Ends of infinite degree correspond to a well-connected sub-structure of a graph as well. The classification provided by Theorem 36 can thus be interpreted to say that an end does not define a well-connected object whenever it defines a non-closed k-tangle for some integer k.

To establish Theorem 36 we delve deeper into the connections between the parameters of the end in question and the topological properties of the k-tangles it defines. Somewhat surprisingly in doing so we can draw a further connection reaching back to Theorem 30 about (finite) tangles being decided by majority vote on a suitable vertex set. It turns out that the k-tangle defined by some end ω is closed if and only if $\deg(\omega) + \operatorname{dom}(\omega) \geqslant k$, which is the case if and only if we can find a vertex set X of finite size (in fact, of size exactly k) such that for each separation (A, B) in the k-tangle a strict majority of X is contained in B:

Theorem 37. *Let τ be the end tangle induced by an end ω of G and let $k \in \mathbb{N}$. Then the following are equivalent:*

(i) *$\tau \cap \vec{S}_k$ is closed in $\vec{S}_k$;*

(ii) *$\deg(\omega) + \operatorname{dom}(\omega) \geqslant k$;*

(iii) *$\tau \cap \vec{S}_k$ has a finite decider set;*

(iv) *$\tau \cap \vec{S}_k$ has a decider set of size exactly k.*

We then turn our attention back to abstract separation systems. It turns out that the prerequisites for the compactness method outlined above can be formulated for abstract separation systems as well, and not just for those of infinite graphs. Separation systems which arise from families of finite separation systems following some compatibility axioms are called *profinite* – in reference to profinite topological spaces – and generalise separation systems of infinite graphs. These profinite systems were recently introduced in [18], which analysed some of their basic combinatorial and topological properties. Similarly as in graphs the property of being a closed subset of a profinite separation system $\vec{S}$ is in a sense equivalent to being amenable to the compactness method.

With this abstraction of separation systems of infinite graphs in place, in Section 5.2 we lift Theorem 18 to these profinite separation systems in a straightforward fashion:

Theorem 38. *Let $\vec{U} = \varprojlim(\vec{U}_p \mid p \in P)$ be a profinite universe of separations and $\mathcal{B}$ a family of non-empty closed subsets of $\vec{U}$. If $\mathcal{B}$ splinters then there is a nested set $N \subseteq \vec{U}$ containing at least one element from each member of $\mathcal{B}$.*

The most important instance of profinite separations are separations of infinite graphs. It thus comes at no surprise that the main application of Theorem 38 is to infinite graphs as well. Roughly speaking, we are able to show that the k-tangles of a graph G that are not included in any tangle of $\vec{S}_{\aleph_0}(G)$ can be arranged in a tree-like structure. This structure is *efficient* in the sense that it contains for any two tangles of G a separation (A, B) on which differ, and such that $|A \cap B|$ is as small as possible among all such separations of G:

Theorem 39. *Let $\mathcal{P}$ be a set of robust regular bounded profiles in G. Then there is a nested set N of separations of G which efficiently distinguishes all distinguishable profiles in $\mathcal{P}$.*

We conclude Chapter 5 by pointing out a second way in which one can use the structure provided by a profinite separation system to establish a tree-of-tangles theorem. This second approach is not found outside of this work. For Theorem 38 we assumed that the sets of separations on which the tangles pairwise differ are closed subsets of a profinite $\vec{S}$; but we could also assume that the tangles themselves are closed subsets of $\vec{S}$, similarly to our study of closed end tangles in Section 5.1. The result we then obtain is the following:

Theorem 40. *Let $\vec{S} \subseteq \vec{U}$ be submodular and closed. Then there is a closed nested set $\vec{T} \subseteq \vec{S}$ that distinguishes all closed regular profiles of $\vec{S}$.*

In examining the differences between Theorem 38 and Theorem 40 it turns out that the latter can be obtained as a consequence of the former despite their difference in approach. The reason for this is that for any two closed tangles the set of separations on which they disagree is also a closed set – which is a non-obvious fact, since in arbitrary topological spaces the symmetric difference of two closed sets need not be closed in general.

The above observation is useful not only for theoretical comparisons between our two profinite theorems. It also allows us to sharpen our application to graphs and improve upon Theorem 39 by replacing the condition 'bounded' with 'closed':

Theorem 42. *Let $\mathcal{P}$ be a set of robust regular profiles of G that are closed in $\vec{S}_{\aleph_0}(G)$. Then there is a nested set N of separations of G which efficiently distinguishes all distinguishable profiles in $\mathcal{P}$.*

At first glance it is not obvious that Theorem 42 is indeed a strengthening of Theorem 39. Using Theorem 37 it is easy to construct k-tangles that are topologically closed but not bounded. Conversely we need that all bounded profiles of a graph are also topologically closed. This is indeed the case, but still takes some effort to prove. We do so in Section 5.2.5:

Theorem 43. *Every bounded profile of G is closed.*

Our applications of Theorems 38 and 40 to graphs, Theorem 39 and Theorem 42, can actually be obtained by employing the stronger [4, Theorem 5.12]. However Carmesin's [4] is a large body of work which treats tangles in graphs only, whereas our proofs of these applications are short and rely on the general-purpose theorems Theorem 38 and 40.

Chapter 2

Terms, notation, and tools

Our graph-theoretic terms and notations follow that of [10], unless stated otherwise. In tangle-theoretic definitions and notation we follow [12, 17–19]; we will, however, give an overview of the tangle-theoretic language henceforth.

2.1 Separation systems and universes

A *separation system* $\vec{S} = (\vec{S}, \leqslant, *)$ is a poset $\vec{S}$ together with an order-reversing involution $*$. If we introduce an element of $\vec{S}$ as $\vec{s}$ then we denote its *inverse* under the involution by $\overleftarrow{s} \coloneqq (\vec{s})^*$. The property of the involution $*$ to be order-reversing then means that $\vec{r} \leqslant \vec{s}$ if and only if $\overleftarrow{r} \geqslant \overleftarrow{s}$ for all $\vec{r}, \vec{s} \in \vec{S}$.

We call the elements $\vec{s}$ of $\vec{S}$ *(oriented) separations*. Given such an oriented separation $\vec{s} \in \vec{S}$ we write s for the set $\{\vec{s}, \overleftarrow{s}\}$, which we call the underlying *unoriented separation* of both $\vec{s}$ and its inverse $\overleftarrow{s}$. Conversely, if $s = \{\vec{s}, \overleftarrow{s}\}$ is an unoriented separation, we call $\vec{s}$ and $\overleftarrow{s}$ the two *orientations* of s. If $\vec{S'}$ is a set of oriented separations we denote by S' the underlying unoriented set of $\vec{S'}$, that is, the set $\{s \mid \vec{s} \in \vec{S'}\}$. On the other hand, if S' is a set of unoriented separations, we write $\vec{S'}$ for the set of all oriented separations whose underlying separation lies in S'. Where appropriate, we shall informally use terms that are defined for oriented separation also for unoriented separations and vice-versa.

A separation $\vec{s} \in \vec{S}$ is *small* if $\vec{s} \leqslant \overleftarrow{s}$, and *trivial in* $\vec{S}$ if $\vec{s} < \vec{r}$ and $\vec{s} < \overleftarrow{r}$ for some $r \in S$. In this case we call r the *witness* of the triviality of $\vec{s}$. If $\vec{s}$ is small or trivial in $\vec{S}$ we call $\overleftarrow{s}$ *co-small* or *co-trivial in* $\vec{S}$, respectively. Observe that all trivial separations are small (and hence all co-trivial ones co-small). A separation $\vec{s}$ and its underlying separation s are *degenerate* if $\vec{s} = \overleftarrow{s}$. An unoriented separation s is *nontrivial* in S if neither of its orientations is trivial in $\vec{S}$, and it is *regular* if neither of its orientations is co-small. A set $\vec{S'}$ of oriented separations is *regular* if it contains no co-small separation, in which case we call S' regular as well.

Two separations r and s are *nested* if they have comparable orientations. Two oriented separations $\vec{r}$ and $\vec{s}$ are *nested* if r and s are nested. Note that $\vec{r}$ and $\vec{s}$ being nested does not imply that $\vec{r}$ and $\vec{s}$ are comparable. If r and s are not nested we say they *cross*, and likewise for $\vec{r}$ and $\vec{s}$. A set of oriented or unoriented separations is *nested* if its elements are pairwise nested.

Two oriented separations $\vec{r}$ and $\vec{s}$ are said to *point towards each other* if $\vec{r} \leqslant \overleftarrow{s}$, and to *point away from each other* if $\overleftarrow{r} \leqslant \vec{s}$. A *star* is a set σ of oriented separations whose elements are nondegenerate and pairwise point towards each other.

A *tree set* is a nested set $\tau \subseteq \vec{S}$ none of whose elements is trivial in τ. If $T = (V, E)$ is a graph-theoretical tree, and $\vec{E} = \{(v, w) \mid \{v, w\} \in E\}$ the set of oriented edges of T, we can define a partial order on $\vec{E}$ by letting $(v, w) < (x, y)$ if $\{v, w\} \neq \{x, y\}$ and the unique $\{v, w\}$-$\{x, y\}$-path in T joins w to x. The involution $*$ given by $(v, w) \mapsto (w, v)$ is then order-reversing, making $\tau(T) = (\vec{E}, \leqslant, *)$ a separation system. It is easy to see that $\tau(T)$ is nested and contains no small (and hence no trivial) separations, and we call it the *edge tree set* of T.

A *universe (of separations)* is a separation system $\vec{U} = (\vec{U}, \leqslant, *, \vee, \wedge)$ that comes with pairwise join- and meet operations $\vee$ and $\wedge$ which turn $\vec{U}$ into a lattice. In universes we have the *De Morgan's rule* that $(\vec{r} \vee \vec{s})^* = (\overleftarrow{r} \wedge \overleftarrow{s})$.

If r and s are separations in a universe U of separations, we call a separation $t \in U$ a *corner separation* of r and s if r, s, and t have orientations such that $\vec{r} \vee \vec{s} = \vec{t}$. Observe that if t is a corner separations of r and s then t equals one of r and s if and only if r and s are nested; when working with separations that cross we will often implicitly use the converse assertion that those two crossing separations are distinct from their corner separations. If t is a corner separation of r and s we also call both orientations of t corner separations of $\vec{r}$ and $\vec{s}$.

One of the fundamental tools in tangle theory is the following basic fact, which is sometimes referred to as the 'fish lemma':

Lemma 2.1.1 ([17, Lemma 2.1])**.** *Let U be a universe of separations and $r, s \in U$ two crossing separations. Every $t \in U$ that is nested with both r and s is also nested with all corner separations of r and s.*

Proof. Since r and s do not have comparable orientations, but t is nested with both of them, there are orientations of these three with $\vec{t} \leqslant \vec{r}$ and $\vec{t} \leqslant \vec{s}$. We then have $\vec{t} \leqslant (\vec{r} \vee \vec{s})$ as well as $\vec{t} \leqslant (\vec{r} \vee \overleftarrow{s})$ and $\vec{t} \leqslant (\overleftarrow{r} \vee \vec{s})$ by transitivity. Finally, we have $\vec{t} \leqslant (\vec{r} \wedge \vec{s})$ by the definition of infimum. Thus $\vec{t}$ is comparable with some orientation of each corner separation of r and s. $\square$

A map $f \colon \vec{S} \to \vec{S'}$ between two separation systems *commutes with the involution* if $f(\vec{s})^* = f(\overleftarrow{s})$ for all $\vec{s} \in \vec{S}$, and it is *order-preserving* if $f(\vec{r}) \leqslant f(\vec{s})$ whenever $\vec{r} \leqslant \vec{s}$. (Note that we do not require the converse assertion: $f(\vec{r}) \leqslant f(\vec{s})$ need not imply that $\vec{r} \leqslant \vec{s}$.) A map $f \colon \vec{U} \to \vec{U'}$ between two universes of separations *commutes with joins and meets* if $f(\vec{r} \vee \vec{s}) = f(\vec{r}) \vee f(\vec{s})$ and $f(\vec{r} \wedge \vec{s}) = f(\vec{r}) \wedge f(\vec{s})$ for all $\vec{r}, \vec{s} \in \vec{U}$. A *homomorphism of separation systems* is then a map $f \colon \vec{S} \to \vec{S'}$ between separation systems which is order-preserving and commutes with the involution. Likewise a map $f \colon \vec{U} \to \vec{U'}$ between two universes is a *homomorphism of universes* if it commutes with the involution, joins, and meets. Observe that if f commutes with joins and meets then f is order-preserving, and consequently every homomorphism of universes is also a homomorphism of separation systems.

A homomorphism between separation systems or universes of separations is an *isomorphism* if it is bijective and its inverse map is also a homomorphism of separation systems or universes, respectively. If there is an isomorphism between two separation systems or two universes we call them *isomorphic*.

2.2 Separations of sets and graphs

Let V be a set. The *separation system of separations of* V is denoted by $\mathcal{S} = \mathcal{S}(V)$ and defined on the set of all pairs (A, B) with $A \cup B = V$, where $(A, B)^* := (B, A)$ and $(A, B) \leqslant (C, D)$ if and only if $A \subseteq C$ and $B \supseteq D$. The *universe of separations of* V is then denoted by $\mathcal{U} = \mathcal{U}(V)$ and consists of $\mathcal{S}(V)$ together with pairwise join $(A, B) \vee (C, D) = (A \cup C\,,\, B \cap D)$ and meet $(A, B) \wedge (C, D) = (A \cap C\,,\, B \cup D)$.

There are two special instances of separations of sets, with the most classical being separations of graphs. If $G = (V, E)$ is a graph, the separation system $\mathcal{S}(G)$ of *separations of* G is the subset of those $(A, B) \in \mathcal{S}(V)$ for which G has no edge incident with vertices in both $A \smallsetminus B$ and $B \smallsetminus A$. Observe that if (A, B) and (C, D) are separations of G then both $(A, B) \vee (C, D)$ and $(A, B) \wedge (C, D)$ are separations of G as well, where these joins and meets are to be taken in $\mathcal{U}(V)$. In this way $\mathcal{S}(G)$ becomes the universe $\mathcal{U}(G)$ of separations of G, a sub-universe of $\mathcal{U}(V)$.

The second special type of separations of sets are bipartitions: those separations (A, B) of some set V with disjoint A and B. We write $\mathcal{SB}(V)$ for the separation system of bipartitions of V. Similarly to graph separations this separation system $\mathcal{SB}(V)$ becomes a universe of separations by inheriting join and meet operations from $\mathcal{U}(V)$, and we write $\mathcal{UB}(V)$ for this universe of bipartitions of V.

In places where no confusion is likely we may also write $\vec{S}(V)$ or $\vec{U}(G)$ rather than $\mathcal{S}(V)$ or $\mathcal{U}(G)$, and so on.

2.3 Order functions and submodularity

Let $\vec{U}$ be a universe of separations. A separation system $\vec{S} \subseteq \vec{U}$ is *(structurally) submodular* if for all $\vec{r}$ and $\vec{s}$ in $\vec{S}$ at least one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ also lies in $\vec{S}$.

An *order function* on $\vec{U}$ is a function $|\ |\colon \vec{U} \to \mathbb{R}$ that is non-negative and symmetric, i.e. with $|\vec{s}| = |\overleftarrow{s}| \geqslant 0$ for all $\vec{s} \in \vec{U}$. We also write $|s|$ for $|\vec{s}|$. Such an order function is *submodular* if

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

for all $\vec{r}, \vec{s} \in \vec{U}$.

If an order function $|\ |$ on $\vec{U}$ is given, and k is some real number, we write $\vec{S}_k$ for the set of all $\vec{s}$ in $\vec{U}$ with $|\vec{s}| < k$. If $|\ |$ is submodular then every such separation system $\vec{S}_k$ is structurally submodular.

2.4 Orientations, profiles, and tangles

Let $\vec{S}$ be a separation system. A set $\vec{S'} \subseteq \vec{S}$ is *antisymmetric* if it contains at most one of $\vec{s}$ and $\overleftarrow{s}$ for all nondegenerate $s \in S$. An *orientation* of a set $M \subseteq S$ of unoriented separations is an antisymmetric set $O \subseteq \vec{M}$ containing at least one of $\vec{s}$ and $\overleftarrow{s}$ for each $s \in M$. A *partial orientation* of M is an orientation of some subset of M.

A set O of oriented separations is *consistent* if it contains no $\vec{r}$ and $\vec{s}$ with $r \neq s$ and $\overleftarrow{r} \leqslant \vec{s}$. In particular a consistent set O cannot contain separations that are co-trivial in O. A partial orientation P of S *extends* to a consistent orientation of S if $P \subseteq O$ for some consistent orientation O of S. Clearly, for P to extend to a consistent orientation of S, the elements of P must not be co-trivial in S. The following standard tool asserts that these co-trivial separations are the only obstacles in extending a consistent partial orientation P to one of S:

Lemma 2.4.1 (Extension Lemma [12])**.** *Let $\vec{S}$ be a separation system and P a consistent partial orientation of S.*

(i) *P extends to a consistent orientation of S if and only if no element of P is co-trivial in $\vec{S}$.*

(ii) *Given any maximal element $\vec{p}$ of P, the orientation O in (i) can be chosen with $\vec{p}$ maximal in O if and only if p is nontrivial in S.*

(iii) *If S is nested, then the orientation in (ii) is unique.*

Let $\vec{U}$ be a universe of separations. A set $P \subseteq \vec{U}$ has the *profile property* if

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

Given a set $\mathcal{F} \subseteq 2^{\vec{U}}$ of subsets of $\vec{U}$, a set $P \subseteq \vec{U}$ is said to *avoid* $\mathcal{F}$ if it includes no element of $\mathcal{F}$, that is, if $\mathcal{F} \cap 2^P$ is empty. An *$\mathcal{F}$-tangle* of some set $S \subseteq U$ is a consistent orientation of S that avoids $\mathcal{F}$. A *profile* of $S \subseteq U$ is a consistent orientation of S with property P, or in other words, an $\mathcal{F}$-tangle of S for

$$\mathcal{F} = \left\{\{\vec{r}, \vec{s}, (\vec{r} \vee \vec{s})^*\} \mid \vec{r}, \vec{s} \in \vec{U}\right\}.$$

If an order function $|\;|$ on $\vec{U}$ is given, a *k-profile* of U is a profile of $S_k \subseteq U$. A set $P \subseteq \vec{U}$ is a *profile in* U if it is a k-profile of U for some $k \in \mathbb{R}$.

A separation $s \in U$ *distinguishes* two antisymmetric sets $P, P' \subseteq \vec{U}$ if s is nondegenerate and has an orientation $\vec{s}$ with $\vec{s} \in P$ and $\overleftarrow{s} \in P'$. If in addition $\vec{U}$ has an order function $|\;|$, and $|s|$ is minimal among all separations which distinguish P and P', then s distinguishes them *efficiently*. The sets P and P' are *distinguishable* if some $s \in U$ distinguishes them.

Let $\mathcal{P}$ be a set of partial orientations of U. A set $N \subseteq U$ of separations *distinguishes* $\mathcal{P}$ if each pair of P and P' in $\mathcal{P}$ is distinguished by some $s \in N$. Accordingly N distinguishes $\mathcal{P}$ *efficiently* if for all P and P' in $\mathcal{P}$ some $s \in N$ distinguishes them efficiently.

If $G = (V, E)$ is a graph, then a submodular order function on the universe $\vec{U} = \mathcal{U}(G)$ of separations of G is given by $|(A, B)| = |A \cap B|$. As for abstract separation systems, we write $\vec{S}_k = \vec{S}_k(G)$ for the subset of $\vec{U}$ consisting of all (A, B) with $|(A, B)| < k$. The classical notion of a graph tangle is then that a k-tangle of G is an $\mathcal{F}$-tangle of S_k for

$$\mathcal{F} = \{\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\} \mid G = G[A_1] \cup G[A_2] \cup G[A_3]\}\,.$$

A *tangle* of G is then a k-tangle for some k. A *maximal tangle* of G is a k-tangle of G that is not the subset of some $(k+1)$-tangle of G.

2.5 Tangle-tree duality

We now give the definitions and tools necessary for our treatment of the tangle-tree duality theorem in Section 4.2 and Section 4.3. These definitions are from [19]. However, since a thorough understanding of both the contents and the techniques of [19] is a prerequisite for Sections 4.2 and 4.3 anyway, and the body of terms, notation, and basic lemmas from [19] specific to tangle-tree duality is disproportionally large compared to the definitions necessary for all other sections in this work, we shall refer the reader to [19] rather than repeating everything here. In the remainder of this section we shall state the most important result of [19] and its underlying definitions.

Let U be a universe of separations and $\vec{S} \subseteq \vec{U}$ a separation system. An *S-tree* is a pair (T, α) of a graph-theoretical tree $T = (V, E)$ and a map $\alpha \colon \vec{E} \to \vec{S}$ that satisfies $\alpha(v, w)^* = \alpha(w, v)$ for all oriented edges (v, w) of T in $\vec{E}$. An S-tree (T, α) is *over* a set $\mathcal{F} \subseteq 2^{\vec{S}}$ if $\{\alpha(t', t) \mid (t', t) \in \vec{E}\} \in \mathcal{F}$ for all nodes t of T.

The *tangle-tree duality theorem* and heart of [19] is the following result:

Theorem 1 (Tangle-tree duality theorem [19])**.** *Let U be a finite universe of separations, $S \subseteq U$ 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 $\mathcal{F}$-separable. Then precisely one of the following holds:*

- *there is an S-tree over $\mathcal{F}$;*
- *there is an $\mathcal{F}$-tangle of S.*

The set $\mathcal{F} \subseteq 2^{\vec{S}}$ is *standard for S* if $\{\overleftarrow{s}\} \in \mathcal{F}$ for every $\vec{s} \in \vec{S}$ that is trivial in $\vec{S}$. The system S is $\mathcal{F}$-separable if for all nontrivial and nondegenerate $\vec{r}, \overleftarrow{r'} \in \vec{S}$ with $\vec{r} \leqslant \vec{r'}$ as well as $\{\overleftarrow{r}\} \notin \mathcal{F}$ and $\{\vec{r'}\} \notin \mathcal{F}$ there exists an $s_0 \in S$ with an orientation $\vec{s}_0 \geqslant \vec{r}$ that emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$ and such that $\overleftarrow{s}_0 \geqslant \overleftarrow{r'}$ emulates $\overleftarrow{r'}$ in $\vec{S}$ for $\mathcal{F}$. Here, $\vec{s}_0$ *emulates* such a separation $\vec{r}$ *in $\vec{S}$ for $\mathcal{F}$* if for every star $\sigma = \{\vec{t}_1, \dots, \vec{t}_n\} \in \mathcal{F}$ with $\overleftarrow{r} \notin \sigma$ and $\vec{r} \leqslant \vec{t}_1$ we have

$$f \downarrow^{\vec{r}}_{\vec{s}_0} (\sigma) := \{\vec{t}_1 \vee \vec{s}_0\,, \vec{t}_2 \wedge \overleftarrow{s}_0, \dots, \vec{t}_n \wedge \overleftarrow{s}_0\} \in \mathcal{F}\,.$$

In particular $f \downarrow^{\vec{r}}_{\vec{s}_0} (\sigma)$ must be included in $\vec{S}$. Note that by the assumptions on $\vec{r}$ the element $\vec{t}_1$ of σ is uniquely determined by the relation $\vec{r} \leqslant \vec{t}_1$.

In Section 4.2 we shall use the following slight re-formulation of a lemma from [19], which holds with the same proof:

Lemma 2.5.1 ([19])**.** *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 nontrivial 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 $\alpha(t) \in \mathcal{F}$.*

Chapter 3

Properties of separation systems

Before we venture forth into the world of tangles, let us arm ourselves with a solid foundation of knowledge about the various structures and scenarios we might encounter in our upcoming theorems and their applications. Tangles, by their classical definition from [41], are orientations of separations occurring in graphs. The fundamental insight of modern tangle theory, however, is that most information about the graph underlying those separations is superfluous for the study of its tangles: the structural information contained in the separations themselves – their partial order, their inverses, and their order function – is in most cases already sufficient.

Furthermore, the theory of abstract separation systems and their tangles also encompasses tangles of matroids, tangles defined on cuts of weighted or unweighted (hyper-) graphs, and tangles of even more general structures. While treating separations at this abstract level can make the behaviour of their concrete instances more transparent, we need those concrete types of separation to guide our intuition also when we study abstract separation systems. We are therefore led to consider the representation problem familiar from other algebraic contexts: *Which abstract separation systems can be represented as separations of graphs?* Or as separations of sets such as set bipartitions?

In this chapter we seek to answer these questions by giving combinatorial characterizations of separation systems of graphs and sets, as well as characterizing those separation systems that come from bipartitions of a set – an important special case of set separations. Additionally, we give examples of separation systems which are fundamentally different from separation systems of sets or graphs.

The structure of this chapter is as follows: in Section 3.1, we introduce the terms and notation for these representations used throughout the chapter, and make precise what it should mean that a given separation system has the form of set separations. This chapter's main results, characterizing separation systems and universes consisting of separations of a set or bipartitions of a set, are given in Section 3.2. Concluding these investigations, we treat the important special case of graph separations in Section 3.3. These three sections are joint work with Nathan Bowler and can also be found in [2].

Following that we engage in three smaller and largely independent subjects of study concerning the general properties of separation systems. First, in Section 3.4, we turn the above question on its head and ask which graphs arise as the 'crossing graphs' of separation systems. We then use our insights to obtain an improved representation theorem for separation systems. Next up, in Section 3.5, we turn to submodular separation systems. This structural property of submodularity lies at the heart of virtually all tree-of-tangles theorems, as we shall see in Section 4.1, but has the downside of requiring an 'ambient' universe in which the separation system at hand lives, even if the trees-of-tangles can be found by working entirely within that separation system. We seek to alleviate this need for an ambient universe by showing that any separation system that is submodular measured in itself can be, if need be, embedded into a universe without altering which corner separations exist. For the conclusion of this chapter we stay on the topic of submodularity: in Section 3.6 we ask whether every submodular separation system possesses a separation the deletion of which leaves that system submodular. This question is open in general, and we answer it affirmatively for those systems that arise as an S_k for some order function. These three sections are as of yet unpublished. They are joint work with Christian Elbracht and Maximilian Teegen, excepting Section 3.5, which is joint work with the latter only.

3.1 Terminology

Let us formally define what it shall mean that a separation system can be implemented by set separations. Given a separation system $\vec{S}$, we say that $\vec{S}$ can be *implemented by set separations* if there are a set V and a sub-system $\vec{S'}$ of $\mathcal{S}(V)$ such that $\vec{S}$ and $\vec{S'}$ are isomorphic. Similarly, we say that $\vec{S}$ can be *implemented by bipartitions* (of a set) if there are a set V and a sub-system $\vec{S'}$ of $\mathcal{SB}(V)$ such that $\vec{S}$ and $\vec{S'}$ are isomorphic.

If a separation system $\vec{S}$ can be implemented by set separations or by bipartitions, we call both $\vec{S'}$ and the isomorphism $f\colon \vec{S} \to \vec{S'}$ witnessing this an *implementation* of $\vec{S}$ by set separations or by bipartitions, respectively.

Finally, for a universe $\vec{U}$, we say that $\vec{U}$ can be *strongly implemented by set separations* if there are a set V and a sub-universe $\vec{U'}$ of $\mathcal{U}(V)$ such that $\vec{U}$ and $\vec{U'}$ are isomorphic universes. Similarly, we say that $\vec{U}$ can be *strongly implemented by bipartitions* if there are a set V and a sub-universe $\vec{U'}$ of $\mathcal{UB}(V)$ such that $\vec{U}$ and $\vec{U'}$ are isomorphic.

If a universe $\vec{U}$ can be strongly implemented by set separations or by bipartitions, we call both $\vec{U'}$ and the isomorphism $f\colon \vec{U} \to \vec{U'}$ witnessing this a *strong implementation* of $\vec{U}$ by set separations or by bipartitions, respectively.

Note that, to show that a separation system $\vec{S}$ can be implemented by set separations or by bipartitions, it suffices to find a ground-set V and an injective homomorphism f from $\vec{S}$ to $\mathcal{S}(V)$ or to $\mathcal{SB}(V)$ which is an isomorphism between $\vec{S}$ and its image $f(\vec{S})$. In Section 3.2 most of the proofs will take this approach.

3.2 Set separations and bipartitions of sets

In this section we shall characterize those separation systems that can be implemented by separations of sets or by bipartitions of sets. We start with a simple observation regarding the shape of small separations in set separation systems:

Lemma 3.2.1. *For any set V, the small separations in $\mathcal{S}(V)$ and $\mathcal{U}(V)$ are those of the form (A, V).*

Proof. Such separations are clearly small, since $A \subseteq V$. On the other hand, if (A, B) is small then we have $(A, B) \leqslant (B, A)$ and so $A \subseteq B$. But this implies $B = A \cup B = V$. $\square$

By Lemma 3.2.1 the small separations in a separation system of sets with ground-set V have the following property: for every pair $(A, V), (B, V) \in \mathcal{U}(V)$ we have $(A, V) \leqslant (B, V)^* = (V, B)$. We will show below that this property characterizes separation systems of sets, so let us make it formal: a separation system $\vec{S}$ is *scrupulous* if for every pair $\vec{r}, \vec{s} \in \vec{S}$ of small separations we have $\vec{r} \leqslant \overleftarrow{s}$.

Using the above observation we can characterize the set separation systems as follows:

Theorem 2. *A separation system $\vec{S}$ can be implemented by set separations if and only if it is scrupulous.*

Proof. First we check that $\mathcal{S}(V)$ is scrupulous for any set V, from which it follows that any subsystem is scrupulous and so that any $\vec{S}$ which can be implemented by set separations is scrupulous. Let (A, V) and (A', V) be small separations of $\mathcal{S}(V)$. Then $A \subseteq V$ and $A' \subseteq V$, so $(A, V) \leqslant (V, A')$.

Now suppose that $\vec{S}$ is scrupulous. Let V be the set of all non-co-small elements of $\vec{S}$. For every $\vec{s} \in \vec{S}$ let

$$A_{\vec{s}} := \{\vec{x} \in V \mid \vec{x} \not\geqslant \vec{s}\}$$

and

$$i(\vec{s}) := (A_{\vec{s}}, A_{\overleftarrow{s}}).$$

For every $\vec{s} \in \vec{S}$, there cannot be any $\vec{x} \in V \setminus (A_{\vec{s}} \cup A_{\overleftarrow{s}})$, since then we would have both $\vec{x} \geqslant \vec{s}$ and $\vec{x} \geqslant \overleftarrow{s}$, so that $\overleftarrow{x} \leqslant \vec{s} \leqslant \vec{x}$, which contradicts the fact that $\vec{x} \in V$ isn't co-small. Thus $i(\vec{s}) \in \mathcal{S}(V)$ for any $\vec{s} \in \vec{S}$. We shall show that i is an implementation of $\vec{S}$ by set separations. It is clear from the definition that i is a homomorphism of separation systems, so it remains to check that it is an isomorphism onto its image. That is, we must show that $i(\vec{s}) \leqslant i(\vec{t})$ implies that $\vec{s} \leqslant \vec{t}$.

So suppose that $i(\vec{s}) \leqslant i(\vec{t})$, that is, $A_{\vec{s}} \subseteq A_{\vec{t}}$ and $A_{\overleftarrow{t}} \subseteq A_{\overleftarrow{s}}$. Since $\vec{t} \notin A_{\vec{t}}$ we have $\vec{t} \notin A_{\vec{s}}$. If $\vec{t}$ is not small then $\vec{t} \in V$ and it follows that $\vec{s} \leqslant \vec{t}$. Similarly, $\overleftarrow{s} \notin A_{\overleftarrow{s}}$ and thus $\overleftarrow{s} \notin A_{\overleftarrow{t}}$, so if $\vec{s}$ is not small then $\overleftarrow{s} \in V$ and hence $\vec{s} \leqslant \vec{t}$. But we also have $\vec{s} \leqslant \vec{t}$ in the remaining case that $\vec{s}$ and $\overleftarrow{t}$ are both small, because $\vec{S}$ is scrupulous. $\square$

Not every separation system is scrupulous, as the next example shows:

Example 3.2.2. Let $\vec{S}$ be the separation system consisting of the separations $\{\vec{r}, \overleftarrow{r}\}$ and $\{\vec{s}, \overleftarrow{s}\}$, with the relations $\vec{r} \leqslant \overleftarrow{r}$ as well as $\vec{s} \leqslant \overleftarrow{s}$ and no further (non-reflexive) relations. Then $\vec{r}$ and $\vec{s}$ are small separations with $\vec{r} \not\leqslant \overleftarrow{s}$, so $\vec{S}$ is not scrupulous and hence cannot be implemented by set separations.

Example 3.2.2 demonstrates how any separation system can be modified so as to not have an implementation by set separations: given a scrupulous separation system $\vec{S'}$, one can make this system non-scrupulous by adding a copy of $\vec{S}$ from Example 3.2.2 to $\vec{S'}$, where separations from $\vec{S'}$ are incomparable to those from the copy of $\vec{S}$. The resulting larger separation system will be non-scrupulous and hence have no implementation by set separations.

However, modifying universes of separations to make them non-scrupulous is not as straightforward as for separation systems due to the existence of joins and meets of any two separations. For universes, being scrupulous is equivalent to another condition on the structure of the small separations:

Lemma 3.2.3. *Let $\vec{U}$ be a universe. Then the following are equivalent:*

(i) *$\vec{U}$ is scrupulous, i.e. $\vec{s} \leqslant \overleftarrow{t}$ for all small $\vec{s}, \vec{t} \in \vec{U}$;*

(ii) *$(\vec{s} \vee \vec{t})$ is small for all small $\vec{s}, \vec{t} \in \vec{U}$;*

(iii) *$(\vec{s} \wedge \vec{t})$ is co-small for all co-small $\vec{s}, \vec{t} \in \vec{U}$.*

Proof. To see that (i) implies (ii), let $\vec{s}, \vec{t} \in \vec{U}$ be two small separations with $\vec{s} \leqslant \overleftarrow{t}$ and $\vec{t} \leqslant \overleftarrow{s}$. Since $\vec{s}$ is small we have $\vec{s} \leqslant \overleftarrow{s}$, so $\vec{s} \leqslant (\overleftarrow{s} \wedge \overleftarrow{t})$. Similarly we have $\vec{t} \leqslant \overleftarrow{t}$ by assumption and hence $\vec{t} \leqslant (\overleftarrow{s} \wedge \overleftarrow{t})$. But this implies $(\vec{s} \vee \vec{t}) \leqslant (\overleftarrow{s} \wedge \overleftarrow{t}) = (\vec{s} \vee \vec{t})^*$ and hence (ii).

To see that, conversely, (ii) implies (i), let $\vec{s}, \vec{t} \in \vec{U}$ be two small separations for which $(\vec{s} \vee \vec{t})$ is small. Then

$$\vec{s} \leqslant (\vec{s} \vee \vec{t}) \leqslant (\vec{s} \vee \vec{t})^* = (\overleftarrow{s} \wedge \overleftarrow{t}) \leqslant \overleftarrow{t}.$$

Finally, for the equivalence of (ii) and (iii), note that for all $\vec{s}, \vec{t} \in \vec{U}$ we have $(\vec{s} \vee \vec{t})^* = (\overleftarrow{s} \wedge \overleftarrow{t})$ by De Morgan's law, which immediately implies the desired equivalence. $\square$

Typically, the second condition in Lemma 3.2.3 is slightly easier to work with than the first, and we shall use it in our proof of Theorem 2's analogue for universes.

To prove a characterization of universes which can be strongly implemented by set separations we shall need the following technical lemma, which is more about lattices than about separation systems:

Lemma 3.2.4. *Let L be a distributive lattice and let x, s and t be elements of L with $s \wedge x \leqslant t \wedge x$ and $s \vee x \leqslant t \vee x$. Then $s \leqslant t$.*

Proof. By elementary calculations we have

$$\begin{aligned} s &= s \vee (s \wedge x) \\ &\leqslant s \vee (t \wedge x) \\ &= (s \vee t) \wedge (s \vee x) \\ &\leqslant (t \vee s) \wedge (t \vee x) \\ &= t \vee (s \wedge x) \\ &\leqslant t \vee (t \wedge x) \\ &= t\,. \end{aligned}$$

□

We are now ready to prove an analogue of Theorem 2 for universes. Since every strong implementation of a universe $\vec{U}$ by set separations is also an implementation of $\vec{U}$, viewed as a separation system without joins and meets, every universe which can be strongly implemented using separations of sets must be scrupulous by Theorem 2. However, being scrupulous is not a sufficient condition for a universe to have a strong implementation by set separations: for every set V, the universe $\mathcal{U}(V)$ as well as all sub-universes of it are (easily seen to be) distributive since intersections and unions of sets are distributive.

So, given a distributive and scrupulous universe $\vec{U}$, how can we find a strong implementation of $\vec{U}$? Let us first suppose that $\vec{U}$ already is a sub-universe of $\mathcal{U}(V)$ for some set V, and see whether we can describe V and each $(A,B) \in \vec{U}$ just in terms of $\vec{U}$ itself, without making use of V.

To this end, for each $v \in V$ let A_v be the set of all $(A,B) \in \vec{U}$ with $v \in A$, and let V' be the set of all those A_v. Then we can write any separation $(A,B) \in \vec{U}$ as

$$(A,B) = (\{v \in V \mid (A,B) \in A_v\}, \{v \in V \mid (B,A) \in A_v\})\,.$$

Thus we can define a map $f \colon \vec{U} \to \mathcal{U}(V')$ as

$$(A,B) \mapsto (\{A_v \in V' \mid (A,B) \in A_v\}\,,\, \{A_v \in V' \mid (B,A) \in A_v\}),$$

and it is easy to check that the map f is an isomorphism of universes between $\vec{U}$ and its image in $\mathcal{U}(V')$. Thus, the ground-set V' we defined can be used to obtain a strong implementation of $\vec{U}$ by separations of sets.

In order to mimic this approach in the general case where we do not know already that $\vec{U}$ is a sub-universe of some $\mathcal{U}(V)$, we need to find a collection V' of sets $X \subseteq \vec{U}$ where the sets $X \in V'$ behave similarly to the sets A_v above. To do this, we shall find some combinatorial properties of the sets A_v, and then take V' as the set of all $X \subseteq \vec{U}$ which have those combinatorial properties.

In the scenario above where $\vec{U}$ is a sub-universe of some $\mathcal{U}(V)$, the first notable property of a set $A_v = \{(A,B) \in \vec{U} \mid v \in A\}$ for some $v \in V$ is that A_v is up-closed: if $(A,B) \in A_v$ and $(C,D) \geqslant (A,B)$, then $v \in A$ and $A \subseteq C$, hence $v \in C$ and $(C,D) \in A_v$. Furthermore A_v is closed under taking meets: if $(A,B),(C,D) \in A_v$, then $v \in A$ and $v \in C$, so $v \in A \cap C$ and hence

$$(A,B) \wedge (C,D) = (A \cap C\,,\, B \cup D) \in A_v.$$

Similarly, we get that the complement of A_v in $\vec{U}$ is down-closed and closed under taking joins. Finally, we can say something about the relationship between A_v

and the small separations of $\vec{U}$: namely, that A_v contains the inverse of every small separation of $\vec{U}$. This is because the small separations of $\vec{U}$ have the form (A, V), so $(V, A) \in A_v$ for all of them.

By taking V' as the set of all $X \subseteq \vec{U}$ which have the five properties from the last paragraph, we can prove Theorem 3:

Theorem 3. *A universe of separations $\vec{U}$ can be strongly implemented by set separations if and only if it is distributive and scrupulous.*

Proof. If $\vec{U}$ has a strong implementation then it is scrupulous by Theorem 2 and distributive because $\mathcal{U}(V)$ is distributive for every V.

Now suppose that $\vec{U}$ is distributive and scrupulous. Let V be the set of all $X \subseteq \vec{U}$ such that

(1) X is up-closed in $\vec{U}$ and closed under taking meets;

(2) $\vec{U} \setminus X$ is down-closed in $\vec{U}$ and closed under taking joins;

(3) X contains all co-small elements of $\vec{U}$.

For any $\vec{s} \in \vec{U}$ we have $(\vec{s} \vee \overleftarrow{s})^* = \overleftarrow{s} \wedge \vec{s} \leqslant \vec{s} \vee \overleftarrow{s}$, so $\vec{s} \vee \overleftarrow{s}$ is co-small for all $\vec{s} \in \vec{U}$. Given $X \in V$ and $\vec{s} \in \vec{U}$ we thus have $(\vec{s} \vee \overleftarrow{s}) \in X$. Therefore X must contain at least one of $\vec{s}$ and $\overleftarrow{s}$, as we cannot have $\vec{s}, \overleftarrow{s} \in \vec{U} \setminus X$ by (2).

For any $\vec{s} \in \vec{U}$ let $A_{\vec{s}} := \{X \in V \mid \vec{s} \in X\}$ and $f(\vec{s}) := (A_{\overleftarrow{s}}, A_{\vec{s}})$. By the above argument we have $A_{\vec{s}} \cup A_{\overleftarrow{s}} = V$, so f takes its image in $\mathcal{U}(V)$. This map clearly commutes with the involution, and by (1) and (2) it also commutes with $\wedge$ and $\vee$. It remains to show that f is injective. For this let $\vec{s}, \vec{t} \in \vec{U}$ with $\vec{s} \neq \vec{t}$ be given; we shall show that $f(\vec{s}) \neq f(\vec{t})$. By switching their roles if necessary we may assume that $\vec{s} \not\leqslant \vec{t}$.

Claim 1. *If there is no co-small $\vec{x}_1 \in \vec{U}$ such that $\vec{s} \wedge \vec{x}_1 \leqslant \vec{t}$ then $A_{\vec{s}} \not\subseteq A_{\vec{t}}$.*

To see this, we wish to find a pair (X, Y) of disjoint subsets of $\vec{U}$ such that

(I) $\vec{s} \in X$ and X is up-closed in $\vec{U}$ and closed under taking meets;

(II) $\vec{t} \in Y$ and Y is down-closed in $\vec{U}$ and closed under taking joins;

(III) X contains all co-small elements of $\vec{U}$.

Call such a pair (X, Y) *good*. We will show later that a maximal good pair (X, Y) will then have $X \in V$ and hence witness that $A_{\vec{s}} \not\subseteq A_{\vec{t}}$; let us show first that some good pair exists. To see this, let

$$A := \bigcup_{\overleftarrow{x} \in \mathrm{Small}(\vec{U})} \lfloor (\vec{s} \wedge \vec{x}) \rfloor$$

and $B := \lceil \vec{t} \rceil$, where $\lfloor\ \rfloor$ and $\lceil\ \rceil$ denote the up- and down-closure in $\vec{U}$. Then A satisfies (I) by Lemma 3.2.3 and (III) by construction, and B clearly satisfies (II). Thus (A, B) is a good pair provided A and B are disjoint. Suppose they are not disjoint; then $\vec{t} \in A$ and hence there is some co-small $\vec{x}_1 \in \vec{U}$ with $\vec{t} \in \lfloor (\vec{s} \wedge \vec{x}_1) \rfloor$, which contradicts the premise of Claim 1.

Now let (X, Y) be an inclusion-wise maximal good pair, which exists by Zorn's Lemma. We wish to show $Y = \vec{U} \setminus X$ since that would imply $X \in V$

and in particular $X \in A_{\vec{s}} \setminus A_{\vec{t}}$. Suppose there exists $\vec{r} \in \vec{U} \setminus (X \cup Y)$. By the maximality of X there are $\vec{x}_1 \in X$ and $\vec{y}_1 \in Y$ with $\vec{r} \wedge \vec{x}_1 \leqslant \vec{y}_1$, and by the maximality of Y there are $\vec{x}_2 \in X$ and $\vec{y}_2 \in Y$ with $\vec{x}_2 \leqslant \vec{r} \vee \vec{y}_2$. Set $\vec{x} := \vec{x}_1 \wedge \vec{x}_2$ and $\vec{y} := \vec{y}_1 \vee \vec{y}_2$. Then $\vec{x} \in X$ and $\vec{y} \in Y$ with $\vec{x} \wedge \vec{r} \leqslant \vec{y} \wedge \vec{r}$ and $\vec{x} \vee \vec{r} \leqslant \vec{y} \vee \vec{r}$. Lemma 3.2.4 now implies $\vec{x} \leqslant \vec{y}$, but this contradicts the fact that X is up-closed and disjoint from Y. Therefore $Y = \vec{U} \setminus X$ and $X \in V$, which proves Claim 1.

Claim 2. *If there is no co-small $\vec{x}_2 \in \vec{U}$ such that $\overleftarrow{t} \wedge \vec{x}_2 \leqslant \overleftarrow{s}$ then $A_{\overleftarrow{t}} \not\subseteq A_{\overleftarrow{s}}$.*

Claim 2 follows in exactly the same way as Claim 1 since $\vec{s} \not\leqslant \vec{t}$ is equivalent to $\overleftarrow{t} \not\leqslant \overleftarrow{s}$.

Claim 3. *There cannot be co-small $\vec{x}_1, \vec{x}_2 \in \vec{U}$ such that $\vec{s} \wedge \vec{x}_1 \leqslant \vec{t}$ and $\overleftarrow{t} \wedge \vec{x}_2 \leqslant \overleftarrow{s}$.*

To see this, suppose $\vec{x}_1, \vec{x}_2$ are as in the claim. Let $\vec{x} := \vec{x}_1 \wedge \vec{x}_2$; this is a co-small separation by Lemma 3.2.3. We then have $\vec{s} \wedge \vec{x} \leqslant \vec{t}$ and $\overleftarrow{t} \wedge \vec{x} \leqslant \overleftarrow{s}$. Applying the involution to the latter inequality gives $\vec{s} \leqslant \vec{t} \vee \overleftarrow{x} \leqslant \vec{t} \vee \vec{x}$. Therefore we get that $\vec{s} \wedge \vec{x} \leqslant \vec{t} \wedge \vec{x}$ as well as $\vec{s} \vee \vec{x} \leqslant \vec{t} \vee \vec{x}$, which by Lemma 3.2.4 contradicts the assumption that $\vec{s} \not\leqslant \vec{t}$. This proves Claim 3.

From Claim 3 it follows that the assumption of at least one of Claim 1 or Claim 2 must be satisfied, and hence $f(\vec{s}) \neq f(\vec{t})$, which completes the proof. □

The next example shows that the assumption of distributivity in Theorem 3 is indeed necessary, as there are abstract universes of separations which are not distributive:

Example 3.2.5. Let L be an arbitrary non-distributive lattice. Let $\vec{U}$ be the separation system consisting of an unoriented separation $\{\vec{r}, \overleftarrow{r}\}$ for each $r \in L$, with the following relations: $\vec{r} \leqslant \vec{s}$ and $\overleftarrow{s} \leqslant \overleftarrow{r}$ if and only if $r \leqslant s$ in L, and additionally $\vec{r} \leqslant \overleftarrow{s}$ for all $r, s \in L$. Since L is a lattice every pair of separations in $\vec{U}$ has a meet and a join, so $\vec{U}$ is a universe. Moreover, since L is non-distributive, $\vec{U}$ is non-distributive by construction as well. Therefore there can be no strong implementation of $\vec{U}$ by set separations. Furthermore $\vec{U}$ is scrupulous, showing that the distributivity cannot be omitted from Theorem 4.

Let us now turn to the topic of (strong) implementations by bipartitions. From Lemma 3.2.1 it follows that the only small separation of $\mathcal{SB}(V)$ is $(\emptyset, V)$. This separation is not only small, it is also the least element of $\mathcal{SB}(V)$. So let us call a separation system $\vec{S}$ *fastidious* if we have $\vec{s} \leqslant \vec{t}$ for all small $\vec{s} \in \vec{S}$ and all $\vec{t} \in \vec{S}$. Clearly every fastidious separation system has at most one small separation, and every separation system with an implementation by bipartitions of sets must be fastidious. Furthermore, every fastidious separation system is scrupulous.

Somewhat surprisingly, Theorem 3 directly implies that every distributive and fastidious universe has a strong implementation by bipartitions of sets:

Theorem 4. *A universe of separations $\vec{U}$ can be strongly implemented by bipartitions of sets if and only if it is distributive and fastidious.*

Proof. Suppose first that $\vec{U}$ is a universe which can be strongly implemented by bipartitions of sets. Then $\vec{U}$ is distributive by Theorem 3. Furthermore, since the only small bipartition of a set V is $(\emptyset, V)$, every separation system which can be implemented by bipartitions is fastidious.

For the other direction, suppose that $\vec{U}$ is a distributive and fastidious universe. Then $\vec{U}$ is scrupulous, so by Theorem 3 there is a strong implementation $f\colon \vec{U} \to \vec{U'}$ of $\vec{U}$ by set separations, where $\vec{U'}$ is a sub-universe of $\mathcal{U}(V)$ for some set V. As $(\vec{s} \wedge \overleftarrow{s})$ is small for each $\vec{s} \in \vec{U}$ there exists a small separation $\vec{r} \in \vec{U}$. Since $\vec{U}$ is fastidious this $\vec{r}$ must be the least element of $\vec{U}$. In particular $\vec{r}$ is the unique small element of $\vec{U}$ and we have $\vec{s} \wedge \overleftarrow{s} = \vec{r}$ for all $\vec{s} \in \vec{U}$. By Lemma 3.2.1, the image of $\vec{r}$ under f is of the form $f(\vec{r}) = (X, V)$ for some $X \subseteq V$. Given any $\vec{s} \in \vec{U}$ and $(A, B) \coloneqq f(\vec{s})$ we thus have $(A, B) \wedge (B, A) = (X, V)$, and in particular $A \cap B = X$.

Consider the map $g\colon \vec{U'} \to \mathcal{U}(V \setminus X)$ defined by

$$g(A, B) \coloneqq (A \setminus X \,,\, B \setminus X).$$

Since $X = A \cap B$ for all $(A, B) \in f(\vec{U})$ the map g is a strong implementation of $\vec{U'}$ by bipartitions of $V \setminus X$. The map $h\colon \vec{U} \to \mathcal{U}(V \setminus X)$ defined by $h = g \circ f$ is thus a strong implementation of $\vec{U}$ by bipartitions of $(V \setminus X)$. $\square$

The next example shows that the assumption of distributivity in Theorem 4 cannot be omitted, since there are fastidious universes which are not distributive:

Example 3.2.6. Let $\vec{U}$ be the universe consisting of three unoriented separations $\{\vec{r}, \overleftarrow{r}\}$, $\{\vec{s}, \overleftarrow{s}\}$ and $\{\vec{t}, \overleftarrow{t}\}$ with the relations

$$\vec{r} \leqslant \vec{s} \leqslant \vec{t} \leqslant \overleftarrow{r} \qquad \text{and} \qquad \vec{r} \leqslant \overleftarrow{t} \leqslant \overleftarrow{s} \leqslant \overleftarrow{r}.$$

Then $\vec{U}$ is a fastidious universe with $\vec{r}$ as the least element, but we have

$$\vec{t} = (\vec{s} \vee \overleftarrow{s}) \wedge \vec{t} \neq (\vec{s} \wedge \vec{t}) \vee (\overleftarrow{s} \wedge \vec{t}) = \vec{s},$$

so $\vec{U}$ is not distributive, showing that the assumption of distributivity in Theorem 4 is necessary.

Example 3.2.6 is a modification of the pentagon lattice N_5, which is an elementary example of a non-distributive lattice. The other prototypical non-distributive lattice, the diamond lattice M_3, can be turned into an example of a fastidious non-distributive universe in a similar fashion.

The power of the theory of universes of bipartitions can be seen in [17], where it is applied to obtain an neat proof of the existence of Gomory-Hu trees in finite graphs.

Interestingly, the conclusion of Theorem 2 that every scrupulous separation system can be implemented by sets does not directly imply that every fastidious separation system has an implementation by bipartitions, even though the analogous implication is true for universes as seen in the proof of Theorem 4. The next example illustrates this:

Example 3.2.7. Let V be the three-element set $\{x, y, z\}$ and $\vec{S}$ the separation system containing the three unoriented set separations

$$\{\{x, y\}, \{x, z\}\}, \qquad \{\{x, y\}, \{y, z\}\} \qquad \text{and} \qquad \{\{x, z\}, \{y, z\}\}.$$

Then $\vec{S}$ has no small separations and hence is a fastidious separation system. However since each $v \in V$ lies in $A \cap B$ for some $(A, B) \in \vec{S}$ it is not possible to obtain an implementation of $\vec{S}$ by bipartitions of sets by deleting those elements from V, as we did in the proof of Theorem 4.

In view of Example 3.2.7, if we want to prove that every fastidious separation system can be implemented by bipartitions of sets, we cannot use Theorem 2 but need to find a direct proof. The following example from [13] points us in the right direction:

Example 3.2.8. Let $T = (V, E)$ be a tree and $\vec{S}$ the edge tree set of T. Then $\vec{S}$ has a natural implementation by bipartitions of sets: for each separation $\vec{s} = (v, w) \in \vec{S}$, the removal of the edge $\{v, w\}$ from T partitions the vertices of T into exactly two connected components $C_{(v,w)}$ and $C_{(w,v)}$, which contain v and w respectively. Furthermore, we have $(v, w) \leqslant (x, y)$ for separations in $\vec{S}$ if and only if $C_{(v,w)} \subseteq C_{(x,y)}$. Thus it is easy to check that the map $f \colon \vec{S} \to \mathcal{SB}(V)$ defined by

$$f\left((v, w)\right) \coloneqq \left(C_{(v,w)}\,,\, C_{(w,v)}\right)$$

is an isomorphism between $\vec{S}$ and its image in $\mathcal{SB}(V)$ and hence an implementation of $\vec{S}$ by bipartitions.

Additionally, the vertex set V of T can be described wholly in terms of $\vec{S}$, without referencing the tree T: every vertex v of T induces a unique consistent orientation O_v of $\vec{S}$ by orienting each edge in $E(T)$ towards v.[1] It is easy to see, then, that for $v \in V$ and $(x, y) \in \vec{S}$ we have $v \in C_{(x,y)}$ if and only if $(y, x) \in O_v$.

This observation leads to another way of implementing $\vec{S}$ by bipartitions of a set. Let $\vec{O} = \vec{O}(\vec{S})$ be the set of all consistent orientations of $\vec{S}$, and for $\vec{s} \in \vec{S}$ let $\vec{O}_{\vec{s}}$ be the set of all $O \in \vec{O}$ which contain $\vec{s}$. Let us define the map $g \colon \vec{S} \to \mathcal{SB}(\vec{O})$ by setting

$$g(\vec{s}) \coloneqq \left(\vec{O}_{\overleftarrow{s}}, \vec{O}_{\vec{s}}\right).$$

Using the Extension Lemma 2.4.1 it is straightforward to check that g is an isomorphism between $\vec{S}$ and its image in $\mathcal{SB}(\vec{O})$ and hence an implementation of $\vec{S}$ by bipartitions of a set.

For an arbitrary separation system $\vec{S}$ let $\vec{O} = \vec{O}(\vec{S})$ be the set of consistent orientations of $\vec{S}$. For $\vec{s} \in \vec{S}$ let us write $\vec{O}_{\vec{s}}$ for the set of all orientations $O \in \vec{O}$ which contain $\vec{s}$. For any nondegenerate $\vec{s} \in \vec{S}$ we have $\vec{O} = \vec{O}_{\vec{s}} \,\dot\cup\, \vec{O}_{\overleftarrow{s}}$ as orientations are antisymmetric.

Remarkably, the map g defined in Example 3.2.8 above which maps a separation $\vec{s}$ to $\left(\vec{O}_{\overleftarrow{s}}\,,\,\vec{O}_{\vec{s}}\right)$ is an isomorphism onto its image in $\mathcal{SB}(\vec{O})$ for all regular separation systems $\vec{S}$, not only those which are the edge tree set of some tree as in Example 3.2.8:

Theorem 5 ([13], Theorem 5.2)**.** *Given any regular separation system $\vec{S}$, the map $f \colon \vec{S} \to \mathcal{SB}(\vec{O})$ given by $\vec{s} \mapsto (\vec{O}_{\overleftarrow{s}}, \vec{O}_{\vec{s}})$ is an isomorphism of separation systems between $\vec{S}$ and its image in $\mathcal{SB}(\vec{O})$.*

Due to the context of [13], in [13] Theorem 5 is only formulated for separation systems $\vec{S}$ that are nested (i.e. in which every two unoriented separations can

[1] Additionally, if T is finite, then every consistent orientation of $\vec{S}$ (viewed as orientation of E) points to a unique vertex: each such orientation has precisely one sink.

be oriented so as to be comparable). However, this assumption is not necessary, and indeed, the proof of Theorem 5 given in [13] does not use it.

Since a fastidious separation system has at most one small separation and is thus 'almost regular' we can utilize Theorem 5 to show that every fastidious separation system has an implementation by bipartitions of sets:

Theorem 6. *A separation system $\vec{S}$ can be implemented by bipartitions of sets if and only if it is fastidious.*

Proof. Since the only small bipartition of a set V is $(\emptyset, V)$, any separation system implemented by bipartitions is fastidious.

Now suppose that $\vec{S}$ is a fastidious separation system. If $\vec{S}$ is regular the assertion follows immediately from Theorem 5, so let us suppose that $\vec{S}$ is not regular. Since $\vec{S}$ is fastidious it then has a unique small separation $\vec{s}$.

We consider first the case that s is the only separation in $\vec{S}$, i.e. that $\vec{S} = \{\vec{s}, \overleftarrow{s}\}$. If s is degenerate, so $\vec{s} = \overleftarrow{s}$, then $\vec{S}$ is isomorphic to $\mathcal{SB}(\emptyset)$. And if $\vec{s} \neq \overleftarrow{s}$ then for any non-empty set V the map $f\colon \vec{S} \to \mathcal{SB}(V)$ which maps $\vec{s}$ to $(\emptyset, V)$ and $\overleftarrow{s}$ to $(V, \emptyset)$ is an implementation of $\vec{S}$ by bipartitions of V.

Let us now suppose that $\vec{S}$ has elements other than s. Then s cannot be degenerate: for if $\vec{s} = \overleftarrow{s}$ then, for each $\vec{t} \in \vec{S}$, we would have both $\vec{s} \leqslant \vec{t}$ and $\overleftarrow{s} = \vec{s} \leqslant \overleftarrow{t}$ by assumption. The latter inequality is equivalent to $\vec{t} \leqslant \vec{s}$ and hence implies $\vec{s} = \vec{t}$, contrary to the assumption that s is not the only separation in $\vec{S}$.

Since $\vec{s}$ is the only small separation of $\vec{S}$, the sub-system $\vec{S'}$ obtained from $\vec{S}$ by deleting s is a regular separation system. By applying Theorem 5 to $\vec{S'}$ we obtain an implementation f' of $\vec{S'}$ using $V = \vec{O}(\vec{S'})$ as a ground-set. Let $f\colon \vec{S} \to \mathcal{UB}(V)$ be the extension of f' to $\vec{S}$ which maps $\vec{s}$ to $(\emptyset, V)$ and $\overleftarrow{s}$ to $(V, \emptyset)$. Since $\vec{s} \leqslant \vec{t}$ for each $\vec{t} \in \vec{S}$ and $\vec{s}$ is nondegenerate f clearly is an implementation of $\vec{S}$ by bipartitions of V. $\square$

3.3 Graph separations

Tangles, the study of which motivated the introduction of abstract separation systems in [12], were introduced in [41] in terms of graph separations. Therefore this instance of separation systems is of special interest. In this section we shall characterize those separation systems that can be implemented by graph separations.

Given an undirected graph $G = (V, E)$, we denote with $\mathcal{U}(G)$ the universe of graph separations of G: the sub-universe of $\mathcal{U}(V)$ consisting of all those separations $\{A, B\}$ for which G contains no edge from $A \smallsetminus B$ to $B \smallsetminus A$.[2]

We say that a universe $\vec{U}$ has a *graphic implementation* if there exists a graph G such that $\mathcal{U}(G)$ and $\vec{U}$ are isomorphic. Note that this notion differs slightly from the notions of implementability in Section 3.2: here we ask that $\vec{U}$ is isomorphic to $\mathcal{U}(G)$ itself, and not just to some sub-universe of $\mathcal{U}(G)$. The reason for this difference is that $\mathcal{U}(G) = \mathcal{U}(V)$ for any graph $G = (V, E)$ with no edges, and hence asking for universes which are isomorphic to a sub-universe of $\mathcal{U}(G)$ for some graph G would be the same as asking for those which are isomorphic to a sub-universe of $\mathcal{U}(V)$ for some set V. The latter is our notion of strong implementations that we already discussed in Section 3.2 and Theorem 3.

Furthermore in this section we shall only deal with finite universes and separation systems.

The aim of this section is then to characterize those universes of separations which have a graphic implementation.

Let us start with a couple of simple observations. Any graphic implementation is also a strong implementation, hence every $\vec{U}$ with a graphic implementation must be distributive and scrupulous by Theorem 3. Additionally $\mathcal{U}(G)$ contains each separation of the form (X, V) for $X \subseteq V$, that is to say, $\mathcal{U}(G)$ contains all small separations of $\mathcal{U}(V)$. These have a very particular structure: the set of small separations of $\mathcal{U}(V)$ forms a boolean algebra. In fact we can say a bit more about this algebra: its maximal element will be (V, V), the only degenerate separation in $\mathcal{U}(V)$. Therefore if $\vec{U}$ is to have a graphic implementation then the set of small separations in $\vec{U}$ must form a boolean algebra whose maximal element is degenerate in $\vec{U}$. It is easy to check that this latter condition already implies that $\vec{U}$ is scrupulous.

So, let us show that every finite universe $\vec{U}$ has a graphic implementation provided $\vec{U}$ is distributive and its set of small separations forms a boolean algebra with degenerate maximal element.

Our strategy for finding a graph G whose universe of separations is isomorphic to such a universe $\vec{U}$ will be as follows: first we shall apply Theorem 3 to find some a implementation $f\colon \vec{U} \to \mathcal{U}(V)$ of $\vec{U}$ for some ground-set V. We will take V as the vertex-set of our graph G and need to define the edges of G in such a way that the image of $\vec{U}$ under f in $\mathcal{U}(V)$ is exactly $\mathcal{U}(G)$. That means that for every $\vec{s} \in \vec{U}$ with $f(\vec{s}) = (A, B)$ we cannot join a vertex from $A \smallsetminus B$ to a vertex in $B \smallsetminus A$, as then $f(\vec{s})$ would not be a separation of G. Conversely, if

[2] In fact our definition of a graph separation differs slightly from [41]: if $\{A, B\}$ is a graph separation of some graph G in our sense, then a graph separation as in [41] would additionally bipartition the edges in $G[A \cap B]$ so as to obtain G as the union of two edge-disjoint subgraphs. For both the study of tangles as well as for our purposes this difference in definition is immaterial, and for this reason both this as well as most other recent works on the theory of tangles and separations use the definition given here.

for two vertices $v, w \in V$ there is no $\vec{s} \in \vec{U}$ whose image (A, B) under f has v and w in $A \setminus B$ and $B \setminus A$ respectively then we must make v and w adjacent in G as otherwise there would be a separation of G which lies outside the image of $\vec{U}$ under f. Thus we will define $E(G)$ as the set of all pairs $\{v, w\}$ of vertices for which there is no (A, B) in the image $f(\vec{U})$ with $v \in A \setminus B$ and $w \in B \setminus A$. This way f will take its image in $\mathcal{U}(G)$, but we still need to prove that f is onto on $\mathcal{U}(G)$.

For the proof of surjectivity we will make use of the the structure of the set of small separations in $\vec{U}$. First we shall show that we can choose the strong implementation $f\colon \vec{U} \to \mathcal{U}(V)$ of $\vec{U}$ in such a way that every small separation of $\mathcal{U}(V)$ lies in the image of $\vec{U}$ under f. Then we shall use this to show that for every non-edge vw of G the separation $(V \setminus \{v\}\,,\, V \setminus \{w\})$ lies in the image of $\vec{U}$ under f. The surjectivity of f onto $\mathcal{U}(G)$ will then follow from the two facts that $f(\vec{U})$ is closed under taking joins and meets and that every separation $(A, B) \in \mathcal{U}(G)$ can be obtained from the separations of the above form through taking joins and meets.

Given a finite universe $\vec{U}$ a separation $\vec{s} \in \vec{U}$ is *atomic* if $\vec{s}$ is not the least element of $\vec{U}$, but the only element less than $\vec{s}$ is the least element of $\vec{U}$. Recall that we write $\mathrm{Small}(\vec{U})$ for the set of small separations of $\vec{U}$.

The next lemma shows that for finite universes we can choose the strong implementation in Theorem 3 in such a way that the image of each atomic element of $\vec{U}$ is of the form $(\{v\}\,,\, V)$:

Lemma 3.3.1. *Let $\vec{U}$ be a finite universe. If $\vec{U}$ is distributive and scrupulous then $\vec{U}$ has a strong implementation by sets in which the image of every small atomic element of $\vec{U}$ is of the form $(\{v\}, V)$ for some v in the ground-set V.*

Proof. Let V be defined as in the proof of Theorem 3: as the set of all $X \subseteq \vec{U}$ with the properties that X is up-closed and closed under taking meets, that $\vec{U} \setminus X$ is down-closed and closed under taking joins, and that X contains all co-small separations. For $\vec{r} \in \vec{U}$ let $A_{\vec{r}}$ be the set of all $X \in V$ with $\vec{r} \in X$. As the proof of Theorem 3 goes on to show, the map $\vec{r} \mapsto (A_{\vec{r}}\,,\, A_{\overleftarrow{r}})$ is an isomorphism between $\vec{U}$ and its image in $\mathcal{U}(V)$. We shall show that a slight modification of this already is an implementation of $\vec{U}$ with the desired property that the image of each small atomic separation is of the form $(\{v\}, V)$.

For this, observe first that $\vec{U}$ (as a set) is an element of the ground-set V by definition of V. We have $\vec{U} \in A_{\vec{r}}$ and $\vec{U} \in A_{\overleftarrow{r}}$ for all $\vec{r} \in \vec{U}$ since $\vec{r}, \overleftarrow{r} \in \vec{U}$. Let $V' := V \setminus \{\vec{U}\}$, and for $\vec{r} \in \vec{U}$ let $A'_{\vec{r}}$ be the set of all $X \in V'$ that contain $\vec{r}$ (equivalently: $A'_{\vec{r}} := A_{\vec{r}} \setminus \{\vec{U}\}$). Then the map $f\colon \vec{U} \to \mathcal{U}(V')$ with $\vec{r} \mapsto (A'_{\vec{r}}\,,\, A'_{\overleftarrow{r}})$ is an isomorphism between $\vec{U}$ and its image in $\mathcal{U}(V')$. We claim that this map f is the desired strong implementation of $\vec{U}$. For this we just need to show that the image under f of each small atomic element of $\vec{U}$ is of the form $(\{v\}, V')$ for some $v \in V'$. So let $\vec{s} \in \vec{U}$ be a small atomic element of $\vec{U}$.

Since $\vec{s}$ is small its image in $\mathcal{U}(V')$ is of the form (Y, V') for some $Y \subseteq V'$. Then Y is non-empty: as $\vec{s}$ is atomic and thus not the least element of $\vec{U}$, its image under f cannot be $(\emptyset, V')$, the least element of $\mathcal{U}(V')$. To finish the proof we thus need to show that Y has at most one element. So let X be an element of $Y = A_{\vec{s}}$. Since $\vec{U}$ is finite and $X \in V'$ is up-closed and closed under taking meets, X is the up-closure of its least element, say $\vec{r}$. As $\vec{s} \in X$ we have $\vec{r} \leqslant \vec{s}$.

But thus, since $\vec{s}$ is atomic, either $\vec{r} = \vec{s}$ or $\vec{r}$ is the least element of $\vec{U}$. The latter would imply $X = \vec{U} \notin V'$, so we must have $X = \lfloor \vec{s} \rfloor$, proving that $A_{\vec{s}}$ contains exactly one element. □

We are now ready to prove that every finite universe that is distributive and whose small elements form a boolean algebra with degenerate maximal element has a graphic implementation:

Theorem 7. *A finite universe $\vec{U}$ has a graphic implementation if and only if it is distributive and* $\mathrm{Small}(\vec{U})$ *is a boolean algebra whose maximal element is degenerate in $\vec{U}$.*

Proof. It is clear that for any graph $G = (V, E)$ the universe $\mathcal{U}(G)$ is distributive and $\mathrm{Small}(\mathcal{U}(G))$ is isomorphic to $\mathcal{P}(V)$, which is a boolean algebra. The maximal small element of $\mathcal{U}(G)$ is (V, V), which is degenerate.

Now suppose that $\vec{U}$ is a distributive universe and that $\mathrm{Small}(\vec{U})$ is a boolean algebra whose maximal element is degenerate. Then in particular there is a maximal small element $\vec{m}$ and for any small elements $\vec{r}, \vec{s} \in \vec{U}$ we have $\vec{r} \leqslant \vec{m} = \overleftarrow{m} \leqslant \overleftarrow{s}$, so that $\vec{U}$ is scrupulous. Thus by Lemma 3.3.1 there are a set V and a map $f \colon \vec{U} \to \mathcal{U}(V)$ such that f is an isomorphism of universes between $\vec{U}$ and its image in $\mathcal{U}(V)$, and for which for every small atomic $\vec{s} \in \vec{U}$ there is a $v \in V$ such that $f(\vec{s}) = (\{v\}, V)$.

Let us now define $G = (V, E)$ to be the graph with vertex set V and an edge from v to w if and only if there is no separation (A, B) in $f(\vec{U})$ with $v \in A \smallsetminus B$ and $w \in B \smallsetminus A$. Note that E is well-defined by this as $(A, B) \in f(\vec{U})$ if and only if $(B, A) \in f(\vec{U})$. Then $\mathcal{U}(G)$ is a sub-universe of $\mathcal{U}(V)$. We claim that f is an isomorphism of universes between $\vec{U}$ and $\mathcal{U}(G)$ and thus a graphic implementation of $\vec{U}$. Since f is an isomorphism between $\vec{U}$ and its image in $\mathcal{U}(V)$ it suffices to show that $f(\vec{U}) = \mathcal{U}(G)$.

To see that $f(\vec{U}) \subseteq \mathcal{U}(G)$, let (A, B) be any separation in $f(\vec{U})$. But then (A, B) is also a separation of G: for all $v \in A \smallsetminus B$ and $w \in B \smallsetminus A$, by definition of E, the separation (A, B) itself witnesses that there is no edge between v and w.

For the converse inclusion $\mathcal{U}(G) \subseteq f(\vec{U})$, observe first that $f(\vec{U})$ is closed under taking inverses, meets, and joins. Let $\vec{m}$ be the maximal element of $\mathrm{Small}(\vec{U})$. Since $\vec{m}$ is degenerate by assumption and (V, V) is the only degenerate separation of $\mathcal{U}(G)$ we must have $f(\vec{m}) = (V, V)$, and since $\mathrm{Small}(\vec{U})$ is a boolean algebra, $\vec{m}$ is the join of small atomic elements. All small atomic separations $\vec{s}$ of $\vec{U}$ get mapped to separations of the type $(\{v\}, V)$ with $v \in V$. The join of those separations can only be $f(\vec{m}) = (V, V)$ if every separation of the form $(\{v\}, V)$ lies in $f(\vec{U})$. Thus we must have $(\{v\}, V) \in f(\vec{U})$ for every $v \in V$, from which it follows that $\mathrm{Small}(\mathcal{U}(V)) \subseteq f(\vec{U})$: for any $X \subseteq V$ we have

$$(X, V) = \bigvee_{x \in X} (\{x\}, V),$$

with the right hand side lying in $f(\vec{U})$. Since $f(\vec{U})$ is closed under taking inverses we also have $(V, X) \in f(\vec{U})$ for every $X \subseteq V$.

Let us now show that for every non-edge vw of G the separation

$$(V \smallsetminus \{v\}, V \smallsetminus \{w\})$$

lies in $f(\overrightarrow{U})$. To see this, let v and w be two non-adjacent vertices of G. The fact v and w are not adjacent means that by definition of E there is some $(A, B) \in f(\overrightarrow{U})$ with $v \in A \setminus B$ and $w \in B \setminus A$. But then $(B, A) \in f(\overrightarrow{U})$, too, and

$$(V \setminus \{v\}\,,\, V \setminus \{w\}) = [(B, A) \vee (V \setminus \{v\}\,,\, V)] \wedge (V\,,\, V \setminus \{w\}),$$

where the right hand side lies in $f(\overrightarrow{U})$.

To finish our proof that $\mathcal{U}(G) \subseteq f(\overrightarrow{U})$, let $(A, B) \in \mathcal{U}(G)$ be arbitrary. Then we can write (A, B) as

$$(A, B) = \bigwedge_{v \in V \setminus A} \bigvee_{w \in V \setminus B} (V \setminus \{v\}\,,\, V \setminus \{w\}),$$

where the right hand side lies in $f(\overrightarrow{U})$ as for all $v \in V \setminus A$ and $w \in V \setminus B$ there can be no edge between v and w as (A, B) is a separation of G, giving $(V \setminus \{v\}\,,\, V \setminus \{w\}) \in f(\overrightarrow{U})$ by the above argument. $\square$

The above theorem characterises those separation systems that can be represented by the set of *all* separations of a (finite) graph. However, in the theory of tangles and separation systems, it is more common to study the separations of a graph up to a given order: the separation system $\vec{S}_k(G)$, which consists of all separations (A, B) of the graph G with $|A \cap B| < k$ for some integer k. Observe that this separation system $\vec{S}_k(G)$ is, in general, not a sub-universe of $\mathcal{U}(G)$. It would be interesting to characterise those separation systems that can be represented as $\vec{S}_k(G)$ for some k and G; unfortunately we were not able to find any nontrivial such characterisation.

It therefore remains an open problem to characterise those separation systems that can be represented by the separations of a graph up to some given order.

3.4 Crossing graphs

Let us now turn the question 'Which separation systems can be represented by the separations of a graph?' on its head and see which graphs can be represented by the crossing graphs of a suitable separation systems. These crossing graphs are a natural concept: given a separation system S we define the *crossing graph* G_S *of* S as the graph $G_S = (V, E)$ with vertex set $V = S$ and $rs \in E$ if and only if r and s cross. (Conversely we could define the nestedness graph of S.)

This concept of crossing graphs is somewhat similar to intersection graphs, interval graphs, and circle graphs, which are all well-studied classes of graphs with unique structural properties. The question we seek to answer in this section is thus: which graphs arise as the crossing graphs of some separation system?

Observe that distinct separation systems of the same size can give rise to the same crossing graph: the crossing graph of a tree set, for instance, has no edges, yet there are plenty of non-isomorphic tree sets.

The answer to our question is, somewhat surprisingly, that each graph arises as the crossing graph of some separation system:

Theorem 8. *For every graph G, not necessarily finite, there is a separation system S such that G is isomorphic to G_S.*

Proof. Let $G = (V, E)$ be a graph. Since the claim is trivial for graphs with fewer than three vertices we may assume that $|V| \geqslant 3$.

Let $V' = V \cup E$. We define a separation system S of bipartitions of V' by letting $\vec{S}$ consist of all (A_v, B_v), where $v \in V$ and

$$A_v = \{v\} \cup E(v)\,, \qquad B_v = V' \setminus A_v\,.$$

We claim that G_S is isomorphic to G. To see this we need to show that any two (A_v, B_v) and (A_w, B_w) in $\vec{S}$ cross if and only if $vw \in E$.

If vw is not an edge in E then the sets A_v and A_w are disjoint, and hence $(A_v, B_v) \leqslant (B_w, A_w)$, showing that these separations are nested.

For the converse, suppose that vw is an edge of G and let us show that (A_v, B_v) and (A_w, B_w) cross. Observe first that A_v and A_w are incomparable. Furthermore A_v is not a subset of B_w since $vw \in A_v \setminus B_w$. Finally, A_v is not a superset of B_w either, since for any $z \in V \setminus \{v, w\}$ we have $z \in B_w$ but $z \notin A_v$. Consequently A_v is incomparable with both A_w and B_w, showing that these separations cross. $\square$

The separation system constructed in Theorem 8 has a curious property: we can find a bijection between the consistent orientations of S and the cliques of G (including the empty set). Indeed, observe that (A_v, B_v) and (A_w, B_w) either cross or point towards each other. Consequently an orientation O of this S is consistent if and only if the set those $v \in V$ for which O contains (B_v, A_v) is a clique in G.

This observation leads to the following more broad claim:

Theorem 9. *If S is a finite regular separation system then there is a bijection between the consistent orientations of S and the cliques of G_S.*

In particular every two regular separation systems with the same crossing graph have the same number of consistent orientations.

Let us first show the following strengthening of the Extension Lemma 2.4.1:

Lemma 3.4.1. *Let S be a regular separation system, $X \subseteq S$ a set of pairwise crossing separations, and P any orientation of X. Then there is a consistent orientation O of S extending P with $P \subseteq (\max O)$.*

Moreover, if P' is a fixed consistent orientation of the set of all those separations in $S \smallsetminus X$ that cross all of X, then there is a unique such O with $P' \subseteq O$.

Proof. We only prove the 'moreover'-part; if no such P' is supplied, we first use the Extension Lemma 2.4.1 to obtain one.

Let $Z \subseteq S$ be the separations in S that are not oriented by $P \cup P'$. By definition of P' every element of Z is nested with some separation in X. Let

$$O \coloneqq P \cup P' \cup \{\vec{s} \mid s \in Z \text{ and } \vec{s} \leqslant \vec{x} \text{ for some } \vec{x} \text{ with } x \in X\} .$$

Then O is antisymmetric and consistent: suppose there are $\vec{r}, \vec{s} \in O$ with $\overleftarrow{r} \leqslant \vec{s}$. Then at least one of them, say $\vec{s}$, lies outside of $P \cup P'$. By definition of O we have $\vec{s} \leqslant \vec{x}$ for some $\vec{x}$ with $x \in X$. Then $\overleftarrow{r} \leqslant \vec{x}$, which means that $\vec{r} \notin (P \cup P')$ since r is nested with x. Therefore there is a separation $\vec{x'}$ with $x' \in X$ and $\vec{r} \leqslant \vec{x'}$. But now x and x' are nested, resulting in a contradiction since $x \neq x'$ by the regularity of S.

Thus O is a consistent orientation of S extending $P \cup P'$. To see that $P \subseteq (\max O)$ consider some $\vec{x} \in P$. Then no other element of $P \cup P'$ lies above $\vec{x}$ since they all cross x. But neither is there a $\vec{s} \in O \smallsetminus (P \cup P')$ with $\vec{x} \leqslant \vec{s}$, since by definition of O this would imply that x is nested with some other element of X. Hence $\vec{x}$ is indeed a maximal element of O.

Finally, for the uniqueness, observe that orienting any $\vec{s} \in O \smallsetminus (P \cup P')$ differently from O would either make O inconsistent (if $\vec{s} \leqslant \vec{x}$ for some $\vec{x} \in P$), or make some element of P not maximal (if $\vec{s} \leqslant \overleftarrow{x}$ for some $\vec{x} \in P$). $\square$

Note that Lemma 3.4.1 holds for infinite separation systems, too.

We can now prove Theorem 9.

Proof of Theorem 9. Let S be a finite regular separation system. Fix a consistent orientation O of S. We will find, for each set $X \subseteq S$ of pairwise crossing separations, a consistent orientation O_X of S such that $(\max O_X) \smallsetminus O$ is an orientation of X.

So let X be a (possibly empty) set of pairwise crossing separations. Let P be the orientation of X given by the inverse of O on X, and let P' be the restriction of O to the set of all those separations in $S \smallsetminus X$ that cross every element of X. By Lemma 3.4.1 there is a unique consistent orientation O_X of S extending $P \cup P'$ with $P \subseteq \max O_X$. Let us show that in fact $P = (\max O_X) \smallsetminus O$. By definition of P and $P \subseteq (\max O_X)$ we have $P \subseteq (\max O_X) \smallsetminus O$. For the converse direction, consider an element $\vec{s} \in (\max O_X) \smallsetminus O$ and suppose that $\vec{s} \notin P$. Then $\vec{s}$ does not lie in P' either since $P' \subseteq O$. Therefore, by the definition of O_X, there is an $\vec{x}$ with $x \in X$ and $\vec{s} \leqslant \vec{x}$. Since $\vec{s}$ was assumed to be a maximal element of O_X we must have $\overleftarrow{x} \in O_X$, i.e. $\overleftarrow{x} \in P$. But then O is inconsistent since O orients s and x as $\overleftarrow{s}$ and $\vec{x}$, a contradiction.

Clearly $O_X \neq O_Y$ if $X \neq Y$. Let us now show that each consistent orientation of S is of the form O_X for some X of pairwise crossing separations.

So let O' be a consistent orientation of S. Let $P \coloneqq (\max O') \smallsetminus O$ and X the set of underlying unoriented separations of P. Then X is a set of pairwise crossing separations: if two separations in P were nested they would point towards each other, which would make O inconsistent since O orients them the other way around. It remains to show that $O' = O_X$. Let Z be the set of all those separations in $S \smallsetminus X$ that cross every element of X. Then O' and O_X agree on X and, by the uniqueness part of Lemma 3.4.1, also on Z. Let P' be the restriction of O_X to $S \smallsetminus (X \cup Z)$. Then $P' \subseteq O$ by definition of O_X. Suppose now that $O_X \neq O'$. Then, since $\vec{S}$ is finite, some maximal element $\vec{s}$ of O' gets oriented as $\overleftarrow{s}$ by O_X. By the above observation we must have $\overleftarrow{s} \in P'$, that is, $\overleftarrow{s} \in O$. But this gives $\vec{s} \in (\max O') \smallsetminus O$ and hence contradicts the fact that O' and O_X agree on X.

Therefore $O' = O_X$, which concludes the proof that the map $X \mapsto O_X$ is a bijection between the pairwise crossing sets and the consistent orientations of S. □

For infinite separation systems the map defined in Theorem 9 is still injective (with the same proof), but not necessarily surjective.

Let us use the technique from Theorem 9 to, in a sense we shall specify below, improve upon Theorem 6, or more accurately, upon Theorem 5. The latter result gives a way to represent a regular separation system S by using as a ground-set V the set of consistent orientations of S, and mapping each $\vec{s} \in \vec{S}$ to the bipartition of V into those orientations which contain $\vec{s}$ and those that do not. Recall that Theorem 5, in its original source [13], was formulated for regular tree sets only. In [32] this result was improved slightly by showing that one can take a much smaller set V: the set of only those consistent orientations with fewer than three maximal elements still gives rise to a representation of any regular tree set, using the same map mapping each $\vec{s}$ to the bipartition of V into the orientations that do and do not contain $\vec{s}$.

As we noted in Section 3.2, Theorem 5 readily extends from regular tree sets to regular separation systems. Let us now show that a similar slimming of the ground-set can be performed there, too:

Theorem 10. *Let S be a regular separation system, O a fixed consistent orientation of S, and O_X defined as in Theorem 9 for every set X of pairwise crossing separations. Let $\mathcal{O}_{<3}$ be the set of all O_X with $|X| < 3$, and for $\vec{s} \in \vec{S}$ let $\mathcal{O}_{<3}(\vec{s})$ be the set of all $O_X \in \mathcal{O}_{<3}$ that contain $\vec{s}$.*

Then the map f given by $\vec{s} \mapsto (\mathcal{O}_{<3}(\overleftarrow{s}), \mathcal{O}_{<3}(\vec{s}))$ is an isomorphism of separation systems between $\vec{S}$ and its image.

Proof. Observe first that $O = O_\emptyset \in \mathcal{O}_{<3}$. The map f clearly commutes with the involution. Furthermore it follows from the consistency of the orientations in $\mathcal{O}_{<3}$ that $f(\vec{r}) \leqslant f(\vec{s})$ whenever $\vec{r} \leqslant \vec{s}$. It thus remains to show that $f(\vec{r})$ and $f(\vec{s})$ cross whenever r and s do, and that f is injective. Once we know the former the latter only needs to be checked for nested r and s.

So let r and s be two crossing separations. Then each of the four possible orientations of r and s is contained in precisely one of $O, O_{\{r\}}, O_{\{s\}}$, and $O_{\{r,s\}}$, which shows that the images of r and s cross as well. In particular the images are distinct.

It remains to show that f is injective. For this we only need to show that $f(r) \neq f(s)$ whenever r and s are distinct nested separations. So let $\vec{r}$ and $\vec{s}$ be orientations of r and s such that $\vec{r} \leqslant \vec{s}$; we need to find an element of $\mathcal{O}_{<3}$ that contains $\vec{r}$ and $\overleftarrow{s}$. If O happens to contain $\vec{r}$ and $\overleftarrow{s}$ there is nothing to show for us. If not, then O contains either $\vec{r}$ and $\vec{s}$, or $\overleftarrow{r}$ and $\overleftarrow{s}$. In the first case $O_{\{r\}}$ is an element of $\mathcal{O}_{<3}$ that contains $\vec{r}$ and $\overleftarrow{s}$, and in the latter case $O_{\{s\}}$ is such an element of $\mathcal{O}_{<3}$. $\square$

In particular if S is finite then the size of the ground-set used to represent S via bipartitions is in $O(|S|^2)$.

Note that Theorem 10 extends easily to fastidious separation systems: similarly to the proof of Theorem 6, the unique small separation of such a fastidious separation system automatically gets mapped to $(\emptyset, V)$.

In general $\mathcal{O}_{<3}$ is not a ground-set of overall smallest possible size to enable a representation of S via bipartitions. This can be seen, for instance, in tree sets: for a tree set of $n-1$ separations the set $\mathcal{O}_{<3}$ has size n, that is, the number of vertices of the tree corresponding to that tree set. However to represent the tree set by bipartitions of some set it suffices to use only the vertices of degree less than three.

3.5 Submodularity

Quite frequently in the remaining chapters we will be studying *structurally submodular* separation systems: separation systems $\vec{S}$ which lie in some ambient universe $\vec{U} \supseteq \vec{S}$, with the structural property that for all pairs $\vec{r}, \vec{s} \in \vec{S}$ at least one of $\vec{r} \vee \vec{s}$ or $\vec{r} \wedge \vec{s}$ again lies in $\vec{S}$, where these joins and meets are taken in $\vec{U}$. However, nearly every time we encounter such a submodular separation system $\vec{S} \subseteq \vec{U}$, we are interested in S only and do not particularly care about the shape of U. The only reason for keeping this ambient universe U around is that we need to be able to express joins and meets of elements of $\vec{S}$, and decide whether these lie in- or outside of $\vec{S}$. The mathematical arguments exploiting the submodularity of S never truly make use of U, but only of the knowledge that at least one of two opposing corner separations is always present in $\vec{S}$.

In this section we offer a way out: a method by which submodularity may be measured and used solely inside $\vec{S}$ itself, without the need for an ambient universe to express this. Let us call a poset $P = (P, \leqslant)$ *submodular* if all $r, s \in P$ have either a pairwise supremum or a pairwise infimum in P. If $\vec{S}$ is a structurally submodular separation system inside some universe $\vec{U}$, then $\vec{S}$ is also submodular as a poset in this sense. We will show a converse to this: if a separation system $\vec{S}$ is submodular as a poset, then we can construct a universe $\vec{U}$ into which $\vec{S}$ embeds in such a way that the pre-existing joins and meets inside $\vec{S}$ are preserved. More precisely, if $\vec{r}$ and $\vec{s}$ have a supremum $\vec{t}$ in $\vec{S}$, then after embedding $\vec{S}$ into $\vec{U}$ we will have $\vec{t} = \vec{r} \vee \vec{s}$, where the latter is measured in $\vec{U}$.

This result allows one to study submodular separation systems in isolation and as independent objects, without introducing the much more involved structure of a universe solely to express a single condition. If, then, at some point during this study the need for such an ambient universe should arise, one can still introduce the universe constructed below, into which the separation system at hand will embed without any changes to its overall of local structure.

Theorem 11. *For every separation system $\vec{S}$ there are a universe $\vec{U}$ and a map $f\colon \vec{S} \to \vec{U}$ that is an isomorphism of separation systems between $\vec{S}$ and its image in $\vec{U}$, with the property that $f(\vec{t}) = f(\vec{r}) \vee f(\vec{s})$ if and only if $\vec{t}$ is the supremum of $\vec{r}$ and $\vec{s}$ in $\vec{S}$, and likewise $f(\vec{u}) = f(\vec{r}) \wedge f(\vec{s})$ if and only if $\vec{u}$ is the infimum of $\vec{r}$ and $\vec{s}$ in $\vec{S}$.*

In particular if $\vec{S}$ is submodular as a poset then $f(\vec{S})$ is submodular as a separation system inside $\vec{U}$.

The heavy lifting of Theorem 11's proof will be done by employing the *Dedekind-MacNeille-completion* [38], a well-known order theoretic tool with which one can embed an arbitrary poset into a suitable lattice while preserving any pre-existing joins and meets. Our task will then be to equip the resulting completion of the poset $\vec{S}$ with an involution which turns it into a universe, and which makes the embedding of $\vec{S}$ into its Dedekind-MacNeille-completion an isomorphism onto its image.

Let us first define this Dedekind-MacNeille-completion. For this we follow the notation of [9]. Let P be any poset, finite or infinite. Given a subset $X \subseteq P$ we write X^ℓ for the set of lower bounds of X in P: the set of all $p \in P$ such that $p \leqslant x$ for all $x \in X$. Similarly we write X^u for the set of all upper bounds of X in P. To improve readability we will omit braces when concatenating these operations, i.e. we shall write $X^{u\ell}$ rather than $(X^u)^\ell$.

The *Dedekind-MacNeille-completion* of P is then given by

$$\mathbf{DM}(P) \coloneqq \left\{X \subseteq P \mid X^{u\ell} = X\right\}$$

using $\subseteq$ as partial order. A nontrivial result by MacNeille [38] asserts that $\mathbf{DM}(P)$ is indeed a lattice. Moreover, the map $f\colon P \to \mathbf{DM}(P)$ given by

$$f(p) \coloneqq \{p\}^{\ell}$$

is an embedding of the poset P into $\mathbf{DM}(P)$ with the property that $f(r)$ is the supremum (resp. infimum) of $f(p)$ and $f(q)$ if and only if r is the supremum (resp. infimum) of p and q in P.

To build some intuition about this Dedekind-MacNeille-completion, observe that for a singleton $\{p\}$, the set $\{p\}^u$ is simply the up-closure $\lfloor p\rfloor$ of p in P. Moreover an element q of P is a lower bound of the up-closure of some p precisely if $q \leqslant p$, and hence $\{p\}^{u\ell} = \{p\}^{\ell} = \lceil p\rceil$. In particular, when applying any series of ^{u} and $^{\ell}$ to a singleton set $\{p\}$, only the very last operation is relevant: for instance $\{p\}^{\ell u\ell} = \{p\}^{\ell}$, which shows that the map f indeed takes its image in $\mathbf{DM}(P)$.

Let us now prove Theorem 11.

Proof of Theorem 11. Let $\vec{S} = (\vec{S}, \leqslant, {}^*)$ be a separation system. Let $\vec{U} = \mathbf{DM}(\vec{S})$ be the Dedekind-MacNeille-completion of $\vec{S}$ with embedding $f\colon \vec{S} \to \vec{U}$ given by $f(\vec{s}) = \{\vec{s}\}^{\ell}$.

For a set $X \subseteq \vec{S}$ we write X^* for the point-wise involution $\{\overleftarrow{x} \mid \vec{x} \in X\}$ of X. For readability we shall extend our convention to omit braces to include *, ^{u}, and $^{\ell}$. Clearly $X^{**} = X$ for all $X \subseteq \vec{S}$.

We define an involution $'$ on $\vec{U}$ by

$$X' \coloneqq X^{u*}$$

and claim that this turns $\vec{U}$ into a universe and f into an isomorphism of separation systems between $\vec{S}$ and its image in $\vec{U}$. To verify this claim we need to ascertain the following: that $'$ takes its image in $\vec{U} = \mathbf{DM}(\vec{S})$; that $'$ is an involution; that $'$ is order-reversing; and finally that f commutes with the involution, i.e. that $f(\vec{s})' = f(\overleftarrow{s})$.

Before we do this, observe that since the involution * of $\vec{S}$ is order-reversing we have

$$X^{u*} = X^{*\ell} \qquad \text{and} \qquad X^{\ell *} = X^{*u}$$

for all $X \subseteq \vec{S}$. We shall be using these two equalities throughout the remainder of the proof.

To see that $'$ takes its image in $\vec{U}$, note that for $X \in \vec{U}$ we have

$$(X')^{u\ell} = X^{u*u\ell} = X^{u\ell*\ell} = X^{*\ell} = X^{u*} = X'\,,$$

where the third equality used the definition of $\vec{U} = \mathbf{DM}(\vec{S})$ to infer $X^{u\ell} = X$. Thus we indeed have $X' \in \vec{U}$ by definition of $\vec{U} = \mathbf{DM}(\vec{S})$.

The map $'$ is an involution since

$$(X')' = X^{u*u*} = X^{u\ell**} = X^{u\ell} = X\,,$$

for $X \in \vec{U}$, using again the definition of $\vec{U} = \mathbf{DM}(\vec{S})$.

To see that $'$ is order-reversing let $X, Y \in \overrightarrow{U}$ with $X \subseteq Y$ be given; we need to show that $X' \supseteq Y'$. From $X \subseteq Y$ it follows that $X^u \supseteq Y^u$, which in turn implies $X^{u*} \supseteq Y^{u*}$. Thus indeed $X' \supseteq Y'$.

To finish the proof it remains to show that $f(\vec{s})' = f(\overleftarrow{s})$ for all $\vec{s} \in \vec{S}$. So let $\vec{s} \in \vec{S}$ be given. Recall that $\{\vec{s}\}^{\ell u} = \{\vec{s}\}^u$. Using this equality we find that

$$f(\vec{s})' = f(\vec{s})^{u*} = \{\vec{s}\}^{\ell u*} = \{\vec{s}\}^{u*} = \{\vec{s}\}^{*\ell} = \{\overleftarrow{s}\}^{\ell} = f(\overleftarrow{s}),$$

as claimed. $\square$

3.6 Unravelling S_k

For the study of structurally submodular separation systems, and for performing proofs by induction, it would be useful if the following assertion were true:

If $S \subseteq U$ is finite and structurally submodular, then so is $S - s$ for some $s \in S$.

If the above assertion were to hold one could 'unravel' every finite submodular separation system, reducing it to the empty set, by deleting one separation at a time without damaging the structural property of submodularity. Proving the above statement is an open problem though.

Curiously it is easy to establish a sort of converse of that statement: rather than looking for an $s \in S$ which one can delete from S, one can always find an $r \in U \setminus S$ which one can *add* to S while keeping it structurally submodular.

Proposition 3.6.1. *If U is finite and $S \subsetneq U$ submodular, then so is $S + r$ for some $r \in U \setminus S$.*

Proof. Let $\vec{r}$ be a maximal element of $\vec{U} \setminus \vec{S}$. Then $S' := S + r$ is submodular: for each $\vec{s} \in \vec{S'}$ we have $(\vec{r} \vee \vec{s}) \in \vec{S'}$ and $(\vec{r} \vee \overleftarrow{s}) \in \vec{S'}$ by the maximality of $\vec{r}$ in $\vec{U} \setminus \vec{S}$. □

Unfortunately inducting on the size of $U \setminus S$ as made possible by Proposition 3.6.1 is rarely useful, and so we shall focus our investigation on unravelling S rather than enlarging it.

A large class of structurally submodular separation systems, and the most important one in practice, arises from a submodular order function on the universe U: if $|\ |$ is a submodular order function then $\vec{S_k}$, the set of all $\vec{s} \in \vec{U}$ with $|\vec{s}| < k$, is structurally submodular. In the remainder of this section we shall demonstrate that the separation systems arising in this way can all be unravelled.

Formally, a submodular separation system S inside some universe U can be *unravelled* if there is an enumeration $S = \{s_1, \dots, s_n\}$ such that $\{s_1, \dots, s_i\}$ is submodular for all $i \leqslant n$. In other words S can be unravelled if we are able to successively delete separations from S until we reach the empty set and maintain structural submodularity throughout.

We show that if U has a submodular order function and $S = S_k$ for some k then S can be unravelled.

Theorem 12. *Let U be a finite universe with a submodular order function $|\ |$ and $S = S_k$ for some k. Then S can be unravelled.*

For the remainder of this section let U be a universe with a submodular order function and $S \subseteq U$. It is easy to see that if $S = S_k$ there is at least one separation which can be deleted without losing submodularity:

Lemma 3.6.2. *If $S = S_k$ and $s \in S$ maximises $|s|$ in S and is nontrivial in $\{r \in S \mid |r| = |s|\}$, then $S - s$ is structurally submodular.*

Proof. Suppose not, that is, suppose that $S - s$ is not structurally submodular. Since S is submodular there are $\vec{r}, \vec{t} \in \vec{S}$ and an orientation $\vec{s}$ of s such that $\vec{r} \vee \vec{t} = \vec{s}$ and $(\vec{r} \wedge \vec{t}) \notin \vec{S}$. By choice of s we have $|s| \geqslant |r|, |t|$ and hence $|\vec{r} \wedge \vec{t}| \leqslant |r|, |t|$ by submodularity. Consequently $(\vec{r} \wedge \vec{t}) \in \vec{S} \setminus s$ unless $(\vec{r} \wedge \vec{t}) = \overleftarrow{s}$. However the latter case is impossible: for if $(\vec{r} \wedge \vec{t}) = \overleftarrow{s}$ then $|s| = |r| = |t|$, and $\vec{s}$ is trivial with both r and t as witnesses, contrary to our assumption on s. Therefore we indeed have $(\vec{r} \wedge \vec{t}) \in \vec{S} \setminus s$, showing that $S - s$ is structurally submodular. $\square$

Unfortunately we cannot rely solely on Lemma 3.6.2 to unravel $S = S_k$, since after its first application and the deletion of s the remaining system $S - s$ may no longer satisfy $S - s = S_{k'}$ for some k'. This happens precisely if s is not the only separation in S of its order, i.e. if there is an $r \in S - s$ with $|r| = |s|$.

To rectify this, and thereby allow the repeated application of Lemma 3.6.2, we shall perturb the order function of U such as to make it injective, whilst maintaining its submodularity and the assertion that $S = S_k$ for a suitable k. For this we show the following:

Theorem 13. *Let U be a finite universe. Then there is a submodular order function $\rho\colon \vec{U} \to \mathbb{N}$ with $\rho(r) \neq \rho(s)$ for all $r \neq s$.*

Proof. Enumerate $\vec{U}$ as $\vec{U} = \{\vec{s}_1, \dots, \vec{s}_n\}$. For $\vec{r} \in \vec{U}$ let $I(\vec{r})$ be the set of all $i \leqslant n$ with $\vec{s}_i \leqslant \vec{r}$. We define $\rho\colon \vec{U} \to \mathbb{N}$ by letting

$$\rho(\vec{r}) = 3^{n+1} - \left(\sum_{i \in I(\vec{r})} 3^i + \sum_{i \in I(\overleftarrow{r})} 3^i \right).$$

This function is clearly symmetric. For the submodularity note that for $\vec{r}$ and $\vec{s}$ in $\vec{U}$ we have $I(\vec{r}) \cap I(\vec{s}) = I(\vec{r} \wedge \vec{s})$ and $I(\vec{r}) \cup I(\vec{s}) \subseteq I(\vec{r} \vee \vec{s})$. Likewise we have $I(\overleftarrow{r}) \cap I(\overleftarrow{s}) = I(\overleftarrow{r} \wedge \overleftarrow{s})$ and $I(\overleftarrow{r}) \cup I(\overleftarrow{s}) \subseteq I(\overleftarrow{r} \vee \overleftarrow{s})$. Therefore each $i \leqslant n$ appears in $I(\vec{r})$ and $I(\vec{s})$ at most as often as it does in $I(\vec{r} \vee \vec{s})$ and $I(\vec{r} \wedge \vec{s})$, and likewise in $I(\overleftarrow{r})$ and $I(\overleftarrow{s})$ at most as often as in $I(\overleftarrow{r} \vee \overleftarrow{s})$ and $I(\overleftarrow{r} \wedge \overleftarrow{s})$. Since $(\vec{r} \vee \vec{s})^* = \overleftarrow{r} \wedge \overleftarrow{s}$ and $(\overleftarrow{r} \vee \overleftarrow{s})^* = \vec{r} \wedge \vec{s}$ this establishes the submodularity.

It remains to show that $\rho(r) \neq \rho(s)$ for all $r \neq s$. For this first note that by definition of ρ we have $\rho(r) = \rho(s)$ if and only if $I(\vec{r}) \cup I(\overleftarrow{r}) = I(\vec{s}) \cup I(\overleftarrow{s})$ and $I(\vec{r}) \cap I(\overleftarrow{r}) = I(\vec{s}) \cap I(\overleftarrow{s})$. It thus suffices to show that $r = s$ whenever $I(\vec{r}) \cup I(\overleftarrow{r}) = I(\vec{s}) \cup I(\overleftarrow{s})$. But this is clear since any maximal element of $\lceil \vec{r} \rceil \cup \lceil \overleftarrow{r} \rceil$ must be either $\vec{r}$ or $\overleftarrow{r}$, where these are the down-closures in $\vec{U}$. $\square$

We can now establish Theorem 12.

Proof of Theorem 12. Let U be a finite universe with submodular order function $|\ |$, which we may assume to be integer-valued. Let $S = S_k$ for some integer k. Let ρ be the order function on U from Theorem 13. Pick a positive constant $c \in \mathbb{R}$ such that $c \cdot \rho(s) < 1$ for all $s \in U$. We define a new order function $\|\ \| : \vec{U} \to \mathbb{R}_{\geqslant 0}$ on U by setting

$$\|s\| \coloneqq |s| + c \cdot \rho(s)\,.$$

Then $\|\ \|$ is submodular and, like ρ, has the property that $\|r\| \neq \|s\|$ whenever $r \neq s$. Enumerate the elements $s_1, \dots, s_n$ of S so that $\|s_1\| < \|s_2\| < \cdots < \|s_n\|$. Then $\{s_1, \dots, s_i\} \subseteq U$ is structurally submodular for each $i \leqslant n$: for $i = n$ it equals S, and for $i < n$ we have that $\{s_1, \dots, s_i\} = S_{\|s_{i+1}\|}$, which is submodular since $\|\ \|$ is. This enumeration therefore demonstrates that S can be unravelled. $\square$

Remark 3.6.3. In Lemma 3.6.2 we showed that if $s \in S = S_k$ maximises $|s|$ in S and is nontrivial in $\{r \in S \mid |r| = |s|\}$ then s can be deleted from S without losing structural submodularity. However it is not clear that deleting s starts an unravelling of S: in the proof of Theorem 12 we only showed that *some* such separation of maximal order in S can be deleted in an unravelling.

It is not difficult to see though that we can indeed start an unravelling of S with any such s: for this we only need to ensure that the function ρ from Theorem 13 has the property that $\rho(s) > \rho(r)$ for all $r \neq s$ with $|r| = |s|$. This can be achieved by taking in the proof of Theorem 13 as the enumeration $\vec{U} = \{\vec{s}_1, \dots, \vec{s}_n\}$ of $\vec{U}$ used to define ρ an enumeration in which $\vec{s}_1, \dots, \vec{s}_\ell$ are those separations with $\vec{s}_i \leqslant \vec{s}$ or $\vec{s}_i \leqslant \overleftarrow{s}$, and $\vec{s}_{\ell+1}, \dots, \vec{s}_n$ all other separations in $\vec{U}$. With this enumeration we have that $\rho(s) \geqslant 3^{n+1} - (3^{\ell+1} - 1)$. If $r \in S - s$ is a separation with $|r| = |s|$, then r has an orientation $\vec{r}$ with $\vec{r} \nleqslant \vec{s}$ and $\vec{r} \nleqslant \overleftarrow{s}$ since s is not trivial with witness r. Then $\vec{r} = \vec{s}_i$ for some $i \geqslant \ell + 1$, and therefore $\rho(r) \leqslant 3^{n+1} - 3^{\ell+1}$, showing that s indeed maximises ρ among all separations of order $|s|$. Consequently s is the element of S which maximises $\| \ \|$ in the proof of Theorem 12, showing that there is an unravelling which deletes s first.

Chapter 4

Finite tangle theory

For our first steps into tangle theory proper we shall concern ourselves with its finite facets; the infinite combinatorics shall have to wait until Chapter 5.

Modern tangle theory rests on two pillars, each of which comprises an archetype of tangle-theorems. The first of these two pillars consists of the *tree-of-tangles theorems*. In these one proves that the tangles of the structure at hand can be separated in a tree-like fashion: there is a nested set of separations such that each pair of tangles differs on some of them. These theorems come in various flavours, depending on their intended setting and application. One can, for instance, ask that each pair of tangles is distinguished as efficiently by the nested set as by the entire host system, where efficiency is measured by some order function. Or one might care about the size or shape of the nested system and its corresponding tree.

In contrast to the first pillar, the second consists of just a single theorem, albeit one that is flexible enough to be applied in a wide variety of applications and for quite a few different purposes: the *tangle-tree duality theorem* Theorem 1. This theorem turns a negative statement into a positive one by asserting that for each sensible set $\mathcal{F}$ of stars, if the host system does not have an $\mathcal{F}$-tangle, then the entire separation system must admit a tree structure which certifies this. The flexibility of this theorem comes from the fact that its user has full control over the set $\mathcal{F}$, which controls not only the characteristics of the $\mathcal{F}$-tangles, but also the shape and properties of the tree resulting from their non-existence. This allows for clever applications even in settings where there are no obvious tangles in sight.

Both of these pillars were, in some form, already present in tangle theory's inaugural work: in [41] Robertson and Seymour established both a tree-of-tangles theorem for tangles in graphs (and hypergraphs), as well as an early precursor of the tangle-tree duality theorem, namely the special case of *brambles* and *tree width* in graphs.

Tangle theory has moved on since then, and more powerful results with more compact and elementary proofs have since been found. In this chapter we shall establish a new state of the art in this respect for both pillars of tangle theory.

In Section 4.1 we formulate the Splinter Theorem, which represents yet another layer of abstraction in the theory of separation systems and tangles, and which easily implies most existing theorems of the first pillar. We then prove it in fifteen lines. We also give a canonical version of the Splinter Theorem, and show that the (non-canonical) Splinter Theorem is general enough the be applied in graph theory proper rather than just in tangle theory. This section up to and including Section 4.1.6 is based on [26], which is joint work with Christian Elbracht and Maximilian Teegen. A version of the splinter theorem not present in [26] is presented in Section 4.1.4. The remaining two parts of Section 4.1, sections 4.1.7 and 4.1.8, are unpublished. The former is joint work with Christian Elbracht and contains a canonical tree-of-tangles theorem not obtainable through the splinter theorem. The latter gives a short proof of a tree-of-tangles theorem using the unravelling technique from Section 3.6.

As for the second pillar, we give two new proofs of Theorem 1 in Section 4.2, each of them somewhat shorter than the original proof from [19]. Both of these new proofs allow for a weakening of the theorem's assumptions. Both of these proofs are my own work. The first proof presented is also given in [24], in which Elbracht, Teegen, and myself make use of the weakening of the assumptions allowed by this new proof in an application. The second proof is unpublished.

Moving on, in Section 4.3 we demonstrate that the two pillars of tangle theory are in reality closer to one-and-a-half pillars: the tangle-tree duality theorem, through tricky choice of $\mathcal{F}$, can be used to yield tree-of-tangle theorems. This novel way of proving results of the first pillar is not just an academic exercise, though: the duality technique allows one to control the degrees in the resulting tree-of-tangles, an aspect never before seen in tree-of-tangles theorems. With the exception of a comparison between the trees-of-tangles found in this way with those constructed in Section 4.1.7, the contents of this section can be found in [24].

Finally, in Section 4.4, we return to the roots of tangle theory and treat classical tangles in graphs. Settling a question raised by Diestel, we show that every graph tangle is *decided* by some multiset of vertices: a (weighted) set of vertices such that, for each low-order separation, the 'big' side according to the tangle is precisely the side containing a majority of those vertices. This result is joint work with Elbracht and Teegen and can be found in [25].

4.1 The tree-of-tangles theorem

4.1.1 Introduction

The central theorem belonging to the first pillar of tangle theory, which was established by Robertson and Seymour together with the notion of tangles, is the following:

Theorem 14 ([41]). *Every graph has a tree-decomposition displaying its maximal tangles.*

Theorem 14 roughly says that the highly cohesive regions in a graph are arranged in a tree-like structure. The 'maximal' in Theorem 14 relates to the order of the tangles: the tree-decomposition found by Theorem 14 displays the graph's tangles at every level of coarseness.

The original proof of Theorem 14 by Robertson and Seymour in [41] is fairly involved and uses as tools multiple nontrivial results about separations in graphs such as, for instance, the existence of certain 'tie-breaker' functions. Since then, the theory of tangles has moved on considerably, and shorter and more elementary proofs of Theorem 14 have been found. The shortest proof to date is due to Carmesin [3, 10], who utilises the fact that the separations needed for the tree-decomposition in Theorem 14 behave well under appropriately defined joins and meets when taken to be of minimal order.

Carmesin, Diestel, Hundertmark, and Stein established the following strengthening of Theorem 14:

Theorem 15 ([6]). *Every graph has a canonical tree-decomposition displaying its maximal tangles.*

Here, 'canonical' means that every automorphism of the graph acts on the decomposition tree. In other words, Theorem 15 uses only invariants of the graph—in particular no tie-breaker—to find the desired tree-decomposition.

A careful analysis of [6]'s proof of Theorem 15 conducted in [17] resulted in another shift of paradigm, similarly to the shift brought about by [41]. Much in the same way as tangles made it possible indirectly to capture substructures in graphs that were traditionally described more directly by sets of vertices or edges, and to treat them in a unified framework, it turned out that tangles themselves could be described, unlike in their definition given by Robertson and Seymour in [41], without reference to vertices or edges.

Indeed, the only information needed about a graph's tangles to prove Theorem 15 is how its separations relate to each other, that is, which separations are nested or cross. Formally, the separations of the graph are turned into a poset together with an involution, and all subsequent tools and theorems in [17] are then formulated for these posets. This new notion of 'abstract tangles' yielded not only a cleaner proof of Theorem 15 in [17], but also made the theory of tangles applicable to a wider range of combinatorial structures.

This novel way of working with so-called 'abstract separation systems' is summarised in [12], and has yielded multiple generalisations and strengthenings of various theorems in the theory of tangles in both the finite (see [2, 13, 14, 16, 17, 19, 20]) and the infinite (see [13, 18, 27, 29, 34]) setting. In this generalised framework tree-decompositions and tangles of graphs are generalised to nested sets of separations and 'profiles', respectively. Working with abstract separation systems rather than with graphs makes many of the results in this new theory of tangles applicable to give notions of highly cohesive substructures in settings other than just graphs, such as in matroids or in image segmentation ([20, 21]).

However one condition not expressed in terms of relations between the separations remained in use throughout the series of abstractions of Theorem 14 implemented in [6] and [17]: all of these works assumed that the separation systems of interest came with a submodular order function. Likewise, Carmesin's short proof of Theorem 14 in [3] also leverages the fact that the order of separations of graphs is a submodular function.

This last non-structural aspect of tree-of-tangles theorems was disposed of in [16]: there Diestel, Erde, and Weißauer replaced the order function with a purely structural notion of submodularity which can be expressed solely in terms of the lattice structure of the separation system. In doing so they established the most general and widely applicable variant of Theorem 14 to date:

Theorem 16 ([16, 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 relevant separation systems in graphs are all structurally submodular, and therefore Theorem 16 still applies to tangles in graphs. On the other hand there are separation systems that are structurally submodular but cannot be represented by graph separations ([2]). In particular Theorem 16 can also be applied to separation systems which, unlike separations in graphs, do not come with any order function, such as arbitrary bipartitions of sets. This is a marked step forward from its predecessor Theorem 14, whose original proof made heavy use of the order of particular separations.

However there is a trade-off involved in Theorem 16's wider applicability: it does not imply Theorem 14. Indeed, Theorem 16 applied to a graph produces a tree-decomposition which displays just the graph's k-tangles for arbitrary but fixed k. This is a significant weakening of Theorem 14, which finds a decomposition displaying the graph's maximal tangles for all tangles orders simultaneously. Moreover, the tree-decomposition found by Theorem 14 is *efficient* in the sense that for every pair of tangles distinguished by the tree-decomposition, the separation in the decomposition distinguishing that pair of tangles is of the lowest possible order. Since Theorem 16 makes only structural assumptions so as to be applicable to separation systems without any order function, Theorem 16 cannot guarantee that the separations used by the nested set to distinguish a particular pair of tangles are of minimal order.

In this section we bridge the gap between Theorem 14 and Theorem 16 by establishing the following tree-of-tangles theorem which combines the upsides of both Theorem 14 and Theorem 16, i.e., which is as widely applicable as Theorem 16 while still being as powerful and efficient as Theorem 14 when applied to tangles in graphs:

Theorem 17. *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 17 includes Theorem 16 by taking a sequence of just one separation system, and it implies Theorem 14 by taking as separation systems S_k the sets of all separations of order $< k$ of the given graph; the resulting nested set is the set of separations of the desired tree-decomposition.

The nested set N found by Theorem 17 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 17, 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:

Theorem 18 (Splinter theorem)**.** *Let U be a universe of separations and $\mathcal{A} = (A_i)_{i \leqslant n}$ a family of subsets of U. If $\mathcal{A}$ splinters then we can pick an element a_i from each A_i so that $\{a_1, \dots, a_n\}$ is nested.*

Theorem 18, 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.

Theorem 18 not only implies Theorem 17, but can also be used to prove Theorem 14 and Theorem 16 directly. In fact Theorem 18 has a remarkably short proof (as we shall see in Section 4.1.2), making it the shortest available proof of Theorem 14 so far (see Section 4.1.3). Moreover, the premise in Theorem 18 is straightforward to check, and Theorem 18 itself does not make reference to tangles or any specific implementations of them. As a result Theorem 18 can be used in many different settings, implying variations of Theorem 14 in a multitude of contexts. For example, after deriving in Section 4.1.3 Theorem 14, Theorem 16, and Theorem 17 from Theorem 18, we use Theorem 18 to establish a new tree-of-tangles theorem in the setting of clique separations.

Since Theorem 18 does not yield a canonical set of separations we cannot deduce Theorem 15 from it. We fix this in Section 4.1.5 by establishing a version of Theorem 18 which does give a canonical nested set, albeit under slightly stronger assumptions:

Theorem 19 (Canonical splinter theorem)**.** *Let U be a universe of separations and let $\mathcal{A} = (A_i \mid i \in I)$ 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 $N = N(\mathcal{A})$ meeting every A_i in $\mathcal{A}$.*

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

We make use of Theorem 19 in Section 4.1.6 to obtain a new shortest proof of Theorem 15 and to extend Theorem 15 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 theorem

In this section we establish our first main theorem of the chapter, Theorem 18, from which we shall derive two previously known results as well as two new flavours of tree-of-tangles theorems in Section 4.1.3. A cornerstone of the proofs of both Theorem 18 as well as of the two known results we shall derive from it is the so-called 'fish lemma' 2.1.1:

Lemma 2.1.1 ([17, Lemma 2.1])**.** *Let U be a universe of separations and $r, s \in U$ two crossing separations. Every $t \in U$ that is nested with both r and s is also nested with all corner separations of r and s.*

Typically, the proof of a tree-of-tangles theorem proceeds by starting with some set N of separations which distinguish some (or all) of the given tangles, and then repeatedly replacing elements r of N which cross some other element s of N with an appropriate corner separation of r and s. Lemma 2.1.1 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 [16].) 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 Theorem 18 seeks to eliminate this careful selection of corner separations for replacement: we will show that for a family $(A_i)_{i \leqslant n}$ of subsets of some universe U we can find a nested set N meeting all the A_i, provided that these sets A_i have one straightforward-to-check property. This theorem will imply many of the existing tree-of-tangles theorems by taking as sets A_i the sets of separations which distinguish the i-th pair of tangles, and checking that the one assumption needed for Theorem 18 is met. Notably, Theorem 18 will make no reference at all to tangles or their specific properties. The proof of Theorem 18 will also utilise Lemma 2.1.1; however, the only assumption we need about the sets A_i is that for elements a_i and a_j of A_i and A_j, respectively, one of their four corner separations lies in either A_i or A_j. This condition will be easy to verify if one wants to deduce other tree-of-tangles theorems from Theorem 18. In fact, the verification of this condition, which just asks for the existence of *some* corner separation of r and s in $A_i \cup A_j$, 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 r and s. So let us define this condition formally.

Let U be a universe and $\mathcal{A} = (A_i)_{i \leqslant n}$ some family of non-empty subsets of U. We say that $\mathcal{A}$ *splinters* if, for every pair $a_i \in A_i$ and $a_j \in A_j$ of separations, either some corner separation of a_i and a_j lies in $A_i \cup A_j$, or one of a_i and a_j lies in $A_i \cap A_j$.

Observe that a family $(A_i)_{i \leqslant n}$ of non-empty sets splinters if and only if for all crossing $a_i \in A_i \setminus A_j$ and $a_j \in A_j \setminus A_i$ one of their four corner separations lies in $A_i \cup A_j$: for if two separations a_i and a_j are nested, these separations themselves are corner separations of the pair a_i and a_j.

With this definition and Lemma 2.1.1 we are already able to state and prove our first main result:

Theorem 18 (Splinter theorem)**.** *Let U be a universe of separations and $\mathcal{A} = (A_i)_{i \leqslant n}$ a family of subsets of U. If $\mathcal{A}$ splinters then we can pick an element a_i from each A_i so that $\{a_1, \dots, a_n\}$ is nested.*

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

Suppose first that we can find some $a_i \in A_i$ such that a_i is nested with at least one element of A_j for each $j \neq i$. Then the assertion holds: for $j \neq i$ let A'_j be the set of those elements of A_j that are nested with a_i. Then $(A'_j \mid j \neq i)$ is a family of non-empty sets which splinters by Lemma 2.1.1. Thus by the induction hypothesis we can pick a nested set $\{a_j \in A'_j \mid j \neq i\}$, 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 $A_1, \dots, A_{n-1}$ to obtain a nested set consisting of some $a_1, \dots, a_{n-1}$. Fix an arbitrary $a_n \in A_n$. For all $i < n$, if a_i itself or one of its corner separations with a_n lies in 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 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 theorem is actually strong enough to directly imply various existing tree-of-tangles theorems, including Theorem 14.

4.1.3 Applications of the splinter theorem

A short proof of Theorem 14

As a first application of Theorem 18 let us give a short proof of Theorem 14:

Theorem 14 ([41])**.** *Every graph has a tree-decomposition displaying its maximal tangles.*

Let us first recall the relevant definitions of separations and tangles in graphs:

Let $G = (V, E)$ be a graph. Then the set $\vec{S} = \vec{S}(G)$ of all separations (A, B) of G is a separation system with involution $(A, B)^* = (B, A)$ and the partial order in which $(A, B) \leqslant (C, D)$ if and only if $A \subseteq C$ and $B \supseteq D$. The *order* of a separation (A, B) of G is $|(A, B)| := |A \cap B|$. With this order function $\vec{S}$ becomes a submodular universe, where $(A, B) \vee (C, D) = (A \cup C\,,\; B \cap D)$. For $(A, B) \in \vec{S}$ we write $\{A, B\}$ for the underlying unoriented separation. Furthermore we write $\vec{S}_k = \vec{S}_k(G)$ for the set of all separations of G of order $< k$.

A k-tangle in G, for an integer k, is an orientation P of $\vec{S}_k$ with the *tangle property:*

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

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 $l > k$.

For a tree-decomposition $(T, \mathcal{V})$ of G and an edge xy of T let T_x and T_y denote the components of $T - xy$ containing x and y, respectively. Let U_x be the union of all bags V_z with $z \in T_x$, and similarly let U_y be the union of all bags V_z with $z \in T_y$. Then we call (U_x, U_y) the *separation induced by xy*. The *set of separations induced by* $(T, \mathcal{V})$ is the set of all separations induced by the edges of T.

We 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.

If N is a nested set of separations of G it is straightforward to find a tree-decomposition of G whose set of induced separations is precisely N (see [13,41]). Therefore, in order to prove Theorem 14, 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

$$A_{P,P'} \coloneqq \{\{A, B\} \in S(G) \mid \{A, B\} \text{ efficiently distinguishes } P \text{ and } P'\}.$$

Since P and P' are not subsets of each other $A_{P,P'}$ is a non-empty set.

Let $\mathcal{A}$ be the family of all these sets $A_{P,P'}$. A nested set of separations of G distinguishes all maximal tangles of G efficiently if and only if it contains an element of each $A_{P,P'}$. Therefore the existence of such a set, and hence Theorem 14, now follows directly from Theorem 18 once we show that $\mathcal{A}$ splinters:

Lemma 4.1.1. *The family $\mathcal{A}$ of all $A_{P,P'}$ splinters.*

Proof. Let $P \neq P'$ and $Q \neq Q'$ be two pairs of distinct maximal tangles of G and let $\{A, B\} \in A_{P,P'}$ and $\{C, D\} \in A_{Q,Q'}$ be two crossing separations. We need to show that we have either $\{A, B\} \in A_{Q,Q'}$ or $\{C, D\} \in A_{P,P'}$, or that some corner separation of $\{A, B\}$ and $\{C, D\}$ lies in $A_{P,P'} \cup A_{Q,Q'}$. 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 $A_{Q,Q'}$. 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 $(A \cup C\,,\, B \cap D)$. 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 $A_{Q,Q'}$.

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 $A_{P,P'}$. □

Abstract tangles in structurally submodular separation systems

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

Theorem 16 ([16, 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 16 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 16 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 14, 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 16 from Theorem 18. 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

$$A_{P,P'} \coloneqq \{s \in S \mid s \text{ distinguishes } P \text{ and } P'\}.$$

For proving Theorem 16 it suffices to show that the family $A_{\mathcal{P}} = (A_{P,P'} \mid P \neq P' \in \mathcal{P})$ splinters:

Lemma 4.1.2. *Given a set $\mathcal{P}$ of profiles of a submodular separation system $\vec{S}$, the family $\mathcal{A}_{\mathcal{P}} = (A_{P,P'} \mid P \neq P' \in \mathcal{P})$ splinters.*

Proof. Let $P \neq P'$ and $Q \neq Q'$ be two pairs of profiles in $\mathcal{P}$ and let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$ be two distinct separations. We need to show that we have either $r \in A_{Q,Q'}$ or $s \in A_{P,P'}$, or that some corner separation of r and s lies in $A_{P,P'} \cup A_{Q,Q'}$. 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 $A_{P,P'}$. $\square$

Let us now deduce Theorem 16 from Theorem 18.

Proof of Theorem 16. Let $\mathcal{P}$ be a set of profiles of S. By Lemma 4.1.2 the collection $(A_{P,P'} \mid P \neq P' \in \mathcal{P})$ of subsets of S splinters. Each of the $A_{P,P'}$ is non-empty as P and P' are distinct profiles of S. Thus, by Theorem 18, we can pick one element from each $A_{P,P'}$ 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 $\mathcal{P}$. $\square$

The above way of using Theorem 18 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 Theorem 18 in the following sections.

Profiles of separations

Theorem 16 from the previous section implied that every graph has, for any fixed integer k, a tree-decomposition which displays its k-tangles. However, Robertson's and Seymour's Theorem 14 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 16 does not imply Theorem 14.

Moreover, since Theorem 16 does not assume that the universe U it is applied to comes with an order function, Theorem 16 cannot say anything about the order of the separations used in the nested set to distinguish all the profiles. If the universe U, as for instance in a graph, *does* come with a submodular order function, one might ask for a nested set which not only distinguishes all the profiles given, but one which does so *efficiently*, i.e., which contains for every pair P, P' of profiles a separation of minimal order among all the separations in U which distinguish P and P'.

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

Theorem 20 (Canonical tangle-tree theorem for separation universes [17, Theorem 3.6])**.** *Let $U = (\vec{U}, \leqslant, {}^*, \vee, \wedge, |\ |)$ be a submodular universe of separations. Then for every robust set $\mathcal{P}$ of profiles in U there is a nested set $T = T(\mathcal{P}) \subseteq U$ of separations such that:*

(i) every two profiles in $\mathcal{P}$ are efficiently distinguished by some separation in T;

(ii) every separation in T efficiently distinguishes a pair of profiles in $\mathcal{P}$;

(iii) for every automorphism α of $\vec{U}$ we have $T(\mathcal{P}^\alpha) = T(\mathcal{P})^\alpha$; (canonicity)

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

Since the definition of robustness of (a set of) profiles is rather involved we do not repeat it here. In the following proofs robustness will be used only in one place; therefore we shall use it there as a black box and refer the reader to [17] for the full definition.

Since every k-tangle of a graph is robust ([17]), Theorem 20 indeed implies Theorem 14 of Robertson and Seymour that every graph has a tree-decomposition displaying its maximal tangles (see [17, Section 4.1] for more on building tree-decompositions from nested sets of separations, and how Theorem 20 implies Theorem 14). Moreover, Theorem 20 improves upon Theorem 14 by finding a *canonical* such tree-decomposition, i.e., one which is preserved by automorphisms of the graph. Since Theorem 18 does not guarantee any kind of canonicity, we are not able to deduce the full Theorem 20 from Theorem 18; however, using Theorem 18 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 20*. (In Section 4.1.5 we shall prove a version of Theorem 18 which implies Theorem 20 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

$$A_{P,P'} := \{a \in U \mid a \text{ distinguishes } P \text{ and } P' \text{ efficiently}\}.$$

Let $\mathcal{A}_\mathcal{P}$ be the family $(A_{P,P'} \mid P \neq P' \in \mathcal{P})$. The only lemma we need in order to apply Theorem 18 is the following:

Lemma 4.1.3. *For a robust set $\mathcal{P}$ of profiles in U the family $\mathcal{A}_\mathcal{P}$ of the sets $A_{P,P'}$ splinters.*

Proof. Let P, P' and Q, Q' be two pairs of distinguishable profiles in $\mathcal{P}$ and let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$ be two crossing separations. We need to show that we have either $r \in A_{Q,Q'}$ or $s \in A_{P,P'}$, or that some corner separation of r and s lies in $A_{P,P'} \cup A_{Q,Q'}$. 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 $r \in A_{Q,Q'}$. 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 $A_{Q,Q'}$. 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. $\square$

Let us now deduce the non-canonical Theorem 20 from Theorem 18:

Proof of the non-canonical Theorem 20. By Lemma 4.1.3 the collection $\mathcal{A}_\mathcal{P}$ of the sets $A_{P,P'}$ splinters. Thus by Theorem 18 we can pick an element from each set $A_{P,P'}$ in $\mathcal{A}_\mathcal{P}$ 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 $A_{P,P'}$, some element of T distinguishes P and P' by definition of $A_{P,P'}$.

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

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

Sequences of submodular separation systems

Let us, once more, compare Theorem 16 and Theorem 20. 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 17 which combines the advantages of both Theorem 16 and Theorem 20 (without canonicity), i.e., which is not dependent on the existence of some order function, but which is as powerful and efficient as Theorem 20 if such an order function does exist.

Concretely, we shall answer the following question, which inspired our main theorem:

If $S_1 \subseteq S_2 \subseteq \ldots \subseteq S_n$ is an ascending sequence of submodular separation 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 $|\ |$ 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 Section 4.1.3 if U has a submodular order function and the S_i are the subsets of U containing all separations of order $< i$.

We also require a structural formulation of the concept of robustness from [17]: 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 $(\vec{r} \vee \vec{s}) \in \vec{S}_j$.

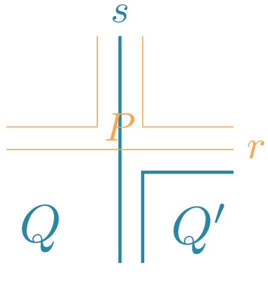

Figure 4.1: Robustness.

With the above definitions we are now able to formally state and prove Theorem 17, which includes both Theorem 16 and the non-canonical Theorem 20 (and hence Theorem 14) as special cases:

Theorem 17. *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 17 runs along very similar lines as the proof of Theorem 20 in the previous section we only sketch it here:

Sketch of proof. For every pair P, P' of distinguishable profiles in $\mathcal{P}$ let $A_{P,P'}$ be the set of all $s \in S_n$ that distinguish P and P' efficiently. The assertion of Theorem 17 follows directly from Theorem 18 if we can show that the family $\mathcal{A}$ of these sets $A_{P,P'}$ splinters.

So let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$ 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 $r \in A_{Q,Q'}$, 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 $A_{Q,Q'}$. 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 $A_{P,P'}$ due to the consistency of P'. □

Theorem 17 directly implies both Theorem 16 and the non-canonical Theorem 20: 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 [31, 46]. More recently clique separations have received quite some attention in theoretical computer science (see for instance [1, 8, 35]) following Tarjan's work [44] on their algorithmic aspects.

In [16] 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 Theorem 18 directly instead of Theorem 16, we are able to obtain a stronger result than the one given in [16], much in the same way that Theorem 17 improves upon Theorem 16.

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. Consequently $\vec{S}_k = \vec{S}_k(G)$ is the set of all clique-separations in G of order less than k, i.e., the set of all $(A, B) \in \vec{S}$ such that $|A \cap B| < k$.

It was shown in [16, 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. $\square$

We can now consider profiles in G with respect to these separation systems. We will use the same notion of profiles as was introduced in [17]: a *profile* P of order k is 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 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

$$A_{P,P'} \coloneqq \{a \in S \mid a \text{ distinguishes } P, P' \text{ efficiently}\}.$$

We will show that the collection of these $A_{P,P'}$ splinters.

Lemma 4.1.5. *For any set $\mathcal{P}$ of profiles the collection $(A_{P,P'} \mid P, P'$ distinguishable profiles$)$ splinters.*

Proof. Let P, P' and Q, Q' be two pairs of distinguishable profiles in $\mathcal{P}$ and let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$ be two distinct separations. We need to show that we have either $r \in A_{Q,Q'}$ or $s \in A_{P,P'}$, or that some corner separation of r and s lies in $A_{P,P'} \cup A_{Q,Q'}$. 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 $r \in A_{Q,Q'}$. 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 $A_{Q,Q'}$. □

It is now straightforward to use Theorem 18 to obtain the following theorem:

Theorem 21. *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 Theorem 18 to $(A_{P,P'} \mid P, P'$ 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 [13] for details). Thus G admits a tree-decomposition whose adhesion sets are cliques and which efficiently distinguishes all the holes and cliques distinguishable by clique separations in G. Such a decomposition is similar to, but not exactly the same as, the decomposition constructed by R. E. Tarjan in [44].

We will see in Section 4.1.6 that such a decomposition can in fact be chosen canonically, i.e., to be invariant under automorphisms of G.

4.1.4 A splinter theorem beyond separations

A careful reading of Section 4.1.2 reveals that throughout both the definitions and the proof of Theorem 18 we only ever worked with *unoriented* separations. Indeed, of the separations considered there the only properties of any relevance to us were whether or not these separations are nested with each other, or corner separations of one another. Both of these notions are defined, of course, as properties arising from the underlying *oriented* separations, but these informations played no part in Theorem 18: to carry out its proof it would have sufficed to know that nestedness is a reflexive symmetric relation between separations, and that corner separations obey the assertion of Lemma 2.1.1. It was only in the applications to profiles that we needed to work with oriented separation.

Let us make this observation concrete by formulating and proving a variant of Theorem 18 which relies on this structural information only, and can be applied in contexts other than separation systems. We start by setting up the necessary definitions in close analogy with Section 4.1.2. Let $\sim$ be a reflexive and symmetric relation on some set S. For $x, y \in S$ we say that x and y are *nested* if $x \sim y$, and otherwise that they *cross*. A subset of S is nested if its elements are pairwise nested, and a single element is nested with some subset of S if it is nested with each element of that set. A *corner* of $x, y \in S$ is an element $z \in S$ with the property that each element of S that crosses z also crosses either x or y.

Note that we have incorporated the assertion of Lemma 2.1.1 for separations into this definition of corner. Observe further that in slight contrast to corners of separations we do not require corners of x and y to be nested with either x or y, and as a consequence both x and y are always corners of themselves. This difference makes the following analogue of 'splinters' slightly more general than its formulation for separations, and also cleaner to state.

A family $\mathcal{A} = (\, A_i\,)_{i \leqslant n}$ of non-empty subsets of S *splinters* if for all $a_i \in A_i$ and $a_j \in A_j$ there is a corner of a_i and a_j which either lies in A_i and is nested with a_j, or else lies in A_j and is nested with a_i.

For such a family $\mathcal{A} = (\, A_i\,)_{i \leqslant n}$ to splinter it suffices that it satisfies the definition only for crossing a_i and a_j with $a_i \in A_i \smallsetminus A_j$ and $a_j \in A_j \smallsetminus A_i$; in all other cases there is nothing to check since a_i and a_j are both corners of a_i and a_j.

The analogue of Theorem 18 then reads as follows:

Theorem 22. *Let $\sim$ be a reflexive and symmetric relation on a set S and $\mathcal{A} = (\, A_i\,)_{i \leqslant n}$ a family of non-empty subsets of S. If $\mathcal{A}$ splinters then we can pick an element a_i from each A_i so that $\{a_1, \dots, a_n\}$ is nested.*

The proof of this relative of Theorem 18, though conceptually unchanged, can be formulated even more succinctly:

Proof. We proceed by induction on n. The assertion clearly holds for $n = 1$, so suppose that it holds for $n - 1$ with $n > 1$. We can thus pick $a_i \in A_i$ for $i < n$ such that $\{a_1, \dots, a_{n-1}\}$ is nested. Choose $a_n \in A_n$ to cross as few of these a_i as possible. If a_n and some a_i cross then any corner of them nested with a_i crosses fewer of the $a_1, \dots, a_{n-1}$ than a_n and hence does not lie in A_n. Thus each A_i contains some element that is nested with a_n; let A'_i be the subset of all those elements. Then $(\, A'_i \mid i < n\,)$ is a family of non-empty sets which splinters. We thus find a nested set $\{a'_1, \dots, a'_{n-1}\}$, which together with a_n is as claimed. □

4.1.5 The canonical splinter theorem

We have seen in Section 4.1.3 that Theorem 18 is already strong enough to imply most of Theorem 20, but crucially does not guarantee the canonicity asserted in (iii). In this section we wish to prove a version of Theorem 18 using a stronger set of assumptions from which we can deduce Theorem 20 in full: we want to find, for a family $\mathcal{A} = (A_i \mid i \in I)$ of subsets of some universe U, a nested set $N = N(\mathcal{A})$ meeting all the A_i that is *canonical*, i.e., which only depends on invariants of $\mathcal{A}$. More formally, we want to find $N = N(\mathcal{A})$ in such a way that if $\mathcal{A}' = (A'_i \mid i \in I)$ 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 $\bigcup_{i \in I} A_i$ and $\bigcup_{i \in I} A'_i$ with $\varphi(A_i) = A'_i$ for all $i \in I$, we ask that $N(\mathcal{A}') = \varphi(N(\mathcal{A}))$. In particular, the nested set found by our theorem should not depend on the universe into which the family $\mathcal{A}$ is embedded.

The assumptions of Theorem 18 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 $\mathcal{A} = (A_1) = (\{s, t\})$. Note that $\mathcal{A}$ splinters, but there may be an automorphism that swaps the two separations, and so the choice of any single one of them would be non-canonical. Since the separations cross 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 [17]. 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 Theorem 18. 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 $\mathcal{A} = (A_i \mid i \in I)$ 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 $\mathcal{A}$ *splinters hierarchically* if for all $a_i \in A_i$ and $a_j \in A_j$ the following two conditions hold:

(1) If $i \prec j$, either some corner separation of a_i and a_j lies in A_j, or two corner separations of a_i and a_j from different sides of a_i lie in A_i.

(2) 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 $c_1 \in A_k$ and $c_2 \in A_i \cup A_j$.

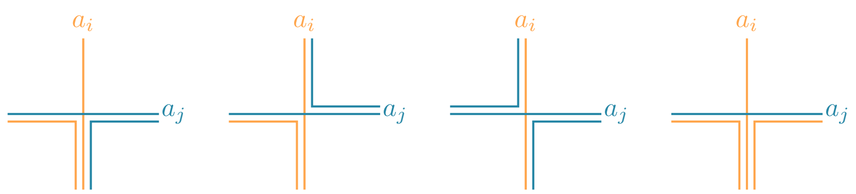

Figure 4.2: The possible configurations in (2) in the definition of *splinter hierarchically*, up to symmetry.

In particular if $\preccurlyeq$ is the trivial partial order on I in which all $i \neq j$ are incomparable then $\mathcal{A}$ splinters hierarchically if and only if (2) holds for all $a_i \in A_i$ and $a_j \in A_j$; 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 A_i in $\mathcal{A}$. The reason we need to allow a partial order $\preccurlyeq$ on I and the slightly weaker condition in (1) for comparable elements of I is that otherwise we would not be able to deduce Theorem 20 in full from our main theorem of this section due to a quirk in the way that robustness is defined for profiles in [17] (see Section 4.1.6).

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

Lemma 4.1.6. *Let $(A_i \mid i \in I)$ 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 $\bigcup_{k\in K} A_k$ is nested.*

Proof. Suppose that $K \subseteq I$ is an antichain and that for some $i, j \in K$ there are $a_i \in A_i$ and $a_j \in A_j$ such that a_i and a_j are extremal in $\bigcup_{k\in K} A_k$ 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 $A_i \cup A_j$, since that would contradict the maximality of $\vec{a}_i$ or $\vec{a}_j$ in $\bigcup_{k\in K} \vec{A}_k$. On the other hand, since neither $i \prec j$ nor $j \prec i$, by condition (2) 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 $A_i \cup A_j$, causing a contradiction to the extremality of a_i and a_j. $\square$

We are now able to prove a canonical splinter theorem by repeatedly applying Lemma 4.1.6 to the collection of the A_i of $\preccurlyeq$-minimal index that have not yet been met by the nested set constructed so far:

Theorem 19 (Canonical splinter theorem)**.** *Let U be a universe of separations and let $\mathcal{A} = (A_i \mid i \in I)$ 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 $N = N(\mathcal{A})$ meeting every A_i in $\mathcal{A}$.*

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

Proof. We proceed by induction on $|I|$. If $|I| = 1$ we can choose as N the set of extremal elements of A_i, which is nested by Lemma 4.1.6 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 4.1.6 the set $E = E(\mathcal{A})$ of extremal elements of $\bigcup_{k\in K} A_k$ is nested. Let $J \subseteq I$ be the set of indices of all those A_j that do not meet E, and for $j \in J$ let A'_j be the set of all elements of A_j that are nested with E. We claim that the collection $\mathcal{A}' = (A'_j \mid j \in J)$ splinters hierarchically with respect to $\preccurlyeq$ on J. This follows from Lemma 2.1.1 as soon as we show that each A'_j is non-empty.

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

Consider first the case that $i \prec j$. By condition (1) of splintering hierarchically, either some corner separation of a_i and a_j lies in A_j, or two corner separations of a_i and a_j from different sides of a_i lie in A_i. The first of these possibilities contradicts the choice of a_j, since that corner separation in A_j would cross fewer elements of E by Lemma 2.1.1. On the other hand, the latter of these possibilities contradicts the choice of a_i as an extremal element of $\bigcup_{k\in K} A_k$. 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 A_j by Lemma 2.1.1. Therefore condition (2) of splintering hierarchically yields the existence of a corner separation of a_i and a_j in A_i for each side of a_i; this, however, contradicts the extremality of a_i in $\bigcup_{k\in K} A_k$ as before.

Therefore each of the sets A'_j with $j \in J$ is non-empty, and hence the collection $\mathcal{A}' = (A'_j \mid j \in J)$ splinters hierarchically with respect to $\preccurlyeq$. Since $|J| < |I|$ we may apply the induction hypothesis to this collection to obtain a canonical nested set $N' = N(\mathcal{A}')$ meeting all A'_j. Now $N = N' \cup E$ is a nested subset of U which meets every 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 $\bigcup_{i\in I} \vec{A_i}$ and a subset of some universe U' such that $\varphi(\mathcal{A})$ splinters hierarchically with respect to $\preccurlyeq$ in U'. Then $\varphi(E) = E(\varphi(\mathcal{A}))$, i.e., the set of extremal elements of $\bigcup_{i\in I} \varphi(A_i)$ is exactly $\varphi(E)$. Therefore $\varphi(E)$ meets $\varphi(A_i)$ if and only if E meets A_i. Consequently the restriction of φ to $\bigcup_{j\in J} \vec{A'_j}$ is an isomorphism of separation systems between $\bigcup_{j\in J} \vec{A'_j}$ and its image in U' with the property that $\varphi(\mathcal{A}')$ splinters hierarchically with respect to $\preccurlyeq$ on J. Moreover, for $j \in J$, the image $\varphi(A'_j)$ of A'_j is exactly the set of those separations in $\varphi(A_j)$ that are nested with $\varphi(E)$.

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

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

concluding the proof. □

4.1.6 Applications of the canonical splinter theorem

In this section we apply Theorem 19 to obtain a short proof of Theorem 20, to strengthen Theorem 21 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 Theorem 19 in the previous section, we are now ready to derive the full version of Theorem 20. For this let $U = (\vec{U}, \leqslant, {}^*, \vee, \wedge, |\ |)$ 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

$$A_{P,P'} := \{a \in U \mid a \text{ distinguishes } P \text{ and } P' \text{ efficiently}\},$$

and let $\mathcal{A}_{\mathcal{P}}$ be the family $(A_{P,P'} \mid \{P, P'\} \in I)$. 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 $A_{P,P'}$ is strictly lower than the order of some element of $A_{Q,Q'}$. Note that the separations in a fixed $A_{P,P'}$ all have the same order.

We shall be able to deduce Theorem 20 from Theorem 19 as soon as we show that $\mathcal{A}_{\mathcal{P}}$ splinters hierarchically.

Lemma 4.1.7. *$\mathcal{A}_{\mathcal{P}}$ splinters hierarchically with respect to $\preccurlyeq$.*

Proof. Let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$ 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 $A_{Q,Q'}$. 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 $A_{Q,Q'}$ and showing that (2) is satisfied.

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

Suppose first that $|r| = |s|$; then neither $\{P, P'\} \prec \{Q, Q'\}$ nor $\{Q, Q'\} \prec \{P, P'\}$. We may assume that P and P' orient s in the same way: for if P and P' orient s differently, we may switch the roles of r and s as well as $\{P, P'\}$ and $\{Q, Q'\}$ and apply the above case. So suppose that both of P and P' contain $\vec{s}$, say. Then neither of the corner separations $\vec{r} \vee \overleftarrow{s}$ nor $\overleftarrow{r} \vee \vec{s}$ can have order strictly less than $|r| = |s|$, as these corner separations would distinguish Q and Q' or P and P', respectively, and would therefore contradict the efficiency of s or of r, respectively. The submodularity of $|\ |$ now implies that both of these corner separations have order exactly $|r| = |s|$ and hence lie in $A_{Q,Q'}$ and $A_{P,P'}$, respectively, showing that (2) 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 $A_{Q,Q'}$, showing that (1) holds and concluding the proof. $\square$

We are now ready to deduce the full Theorem 20 from Theorem 19:

Theorem 20 (Canonical tangle-tree theorem for separation universes [17, Theorem 3.6])**.** *Let $U = (\vec{U}, \leqslant, ^*, \vee, \wedge, |\ |)$ 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 4.1.7 the family $\mathcal{A}_\mathcal{P}$ splinters hierarchically. Thus we can apply Theorem 19 to $\mathcal{A}_\mathcal{P}$ to obtain a nested set $N = N(\mathcal{A}_\mathcal{P})$ which meets every $A_{P,P'}$. Clearly, N satisfies (i), (ii) and (iv) of Theorem 20.

To see that N satisfies (iii), let α be an automorphism of $\vec{U}$. Then the restriction of α to $\bigcup_{\{P,P'\}\in I} \vec{A}_{P,P'}$ is an isomorphism of separation systems onto its image in $\vec{U}$. We therefore have, by Theorem 19, that $\alpha(N(\mathcal{A}_\mathcal{P})) = N(\alpha(\mathcal{A}_\mathcal{P}))$. For every $A_{P,P'}$ in $\mathcal{A}_\mathcal{P}$ we have that $\alpha(A_{P,P'})$ is precisely the set of those separations in U which distinguish P^α and P'^α efficiently; in other words, we have $\alpha(\mathcal{A}_\mathcal{P}) = \mathcal{A}_{\mathcal{P}^\alpha}$, 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 $A_{P,P'}$ splinters, but can be used to show that the collection of these $A_{P,P'}$ even splinters hierarchically, allowing us to apply Theorem 19: 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) $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$.

To see this, let P, P' and Q, Q' be distinguishable pairs of profiles of clique separations. Let $r \in A_{P,P'}$ and $s \in A_{Q,Q'}$, 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 $A_{P,P'}$ and $A_{Q,Q'}$, 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 $A_{Q,Q'}$. Therefore, if $|r| < |s|$, condition (1) of splintering hierarchically is satisfied.

So suppose further that $|r| = |s|$, and let us check that (2) of splintering hierarchically is satisfied. Observe that, similarly as above, P orients s, and P' contains both $(\vec{r} \wedge \overleftarrow{s})$ 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 $A_{P,P'}$. If this corner separation in $A_{P,P'}$ and the corner separation in $A_{Q,Q'}$ found above are from different sides of either r or s, then (2) 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 $A_{P,P'}$, which together with $(\overleftarrow{r} \wedge \overleftarrow{s}) \in A_{Q,Q'}$ witnesses that (2) is fulfilled.

Therefore, by Theorem 19 we get that we can choose the set in Theorem 21 canonically:

Theorem 23. *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 Theorem 19 to the family of the sets $A_{P,P'}$ 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 [16]: given a fixed cyclic order on a ground-set V, a *circle separation* of V is a bipartition (A, B) of V into two disjoint intervals in the cyclic order. Observe that the set of all circle separations is not closed under joins and meets and hence not a sub-universe of the universe of all bipartitions of V:

Example 4.1.8. 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 $U = (\vec{U}, \leqslant, {}^*, \vee, \wedge, |\ |)$ the universe of all bipartitions of V with a submodular order function $|\ |$. Let $S \subseteq U$ be the set of all separations in U that are circle separations of V. Consequently we denote by 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 a consistent orientation of S_k a *k-tangle* in S if it has no subset in

$$\mathcal{F} = \mathcal{F}^n_m := \left\{ F \subseteq 2^{\vec{U}} \;\middle|\; \left|\bigcap_{(A,B)\in F} B\right| < m \text{ and } |F| < n \right\}.$$

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 Theorem 19 we can show that there is a canonical nested set of circle separations which efficiently distinguishes all distinguishable tangles in S:

Theorem 24. *The set S of all circle separations of V contains a tree set $T = T(S)$ that efficiently distinguishes all distinguishable tangles of S. Moreover, this tree set T can be chosen canonically, i.e., so that for every automorphism α of S we have $T(S^\alpha) = T(S)^\alpha$.*

In order to prove Theorem 24 we need the following short lemma:

Lemma 4.1.9. *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 24.

Proof of Theorem 24. For every pair P, P' of distinguishable tangles in S let $A_{P,P'}$ 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 $A_{P,P'}$ have strictly lower order than those in $A_{Q,Q'}$.

Let us show that the collection of these sets $A_{P,P'}$ splinters hierarchically; the claim will then follow from Theorem 19.

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

So suppose that r and s cross. Then by Lemma 4.1.9 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 (1) is satisfied, i.e., to find a corner separation of r and s in $A_{Q,Q'}$. 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 $A_{Q,Q'}$, 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 (2) is satisfied. For this we shall find corner separations from different sides of r or of s that lie in $A_{P,P'}$ and $A_{Q,Q'}$, 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 $A_{Q,Q'}$. Now this corner separation together with $\vec{r} \vee \vec{s}$ witnesses that (2) holds. $\square$

4.1.7 A canonical tree-of-tangles theorem for structural submodularity

Let us complement the previous sections by establishing a tree-of-tangles theorem that is not obtainable through the use of Theorem 18 or Theorem 19. As we have seen there, those two theorems together yield a plethora of tree-of-tangle theorems across a multitude of different settings, incorporating efficiency and canonicity wherever possible. However, there is still one tree-of-tangles theorem with neither an efficient nor a canonical version: Theorem 16, the statement that the profiles of a structurally submodular separation systems can be distinguished by a nested set.

Since the main point of studying separation systems that are structurally submodular is that these still allow for a tree-of-tangles theorem but do no necessarily come from some order function, there is no meaningful notion of 'distinguishing efficiently' in this setting. However one can still ask for a canonical tree-of-tangles theorem, i.e. a canonical way to find a nested set $N(\mathcal{P})$ distinguishing a given set $\mathcal{P}$ of profiles of the submodular S. This cannot be done by using Theorem 19 though: the family $(\, A_{P,P'} \mid P \neq P' \text{ in } \mathcal{P}\,)$ of the sets of separations distinguishing a given pair of profiles does not splinter hierarchically. The reason for this is that if $\vec{r}$ and $\vec{s}$ are separations with $\vec{r}, \vec{s} \in P$ and $\overleftarrow{r}, \overleftarrow{s} \in P'$ for some two profiles in $\mathcal{P}$, then both $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ also distinguish P and P' and hence lie in $A_{P,P'}$ – but only if these corner separations happen to lie in $\vec{S}$. In the presence of an order function and with $\vec{S} = \vec{S}_k$ for some k, and with r and s distinguishing P and P' efficiently, this very efficiency implies that the two corner separations of r and s have exactly the same order as r and s themselves, and hence not only lie in $\vec{S} = \vec{S}_k$ but also in $A_{P,P'}$. Without the help of this order function, though, it is possible that only one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ lies in $\vec{S}$, in which case (2) would be violated.

In the remainder of this section we seek to fill this gap in the splinter theorems' coverage and establish the following canonical version of Theorem 16:

Theorem 25. *Let $\vec{U}$ be a finite universe of separations, $\vec{S} \subseteq \vec{U}$ submodular, and $\mathcal{P}$ a set of profiles of S. Then there is a nested set $N = N(\mathcal{P}) \subseteq S$ which distinguishes $\mathcal{P}$. This $N(\mathcal{P})$ can be chosen canonically: if $\varphi\colon \vec{U} \to \vec{U'}$ is an isomorphism of universes, then $\varphi(N(\mathcal{P})) = N(\varphi(\mathcal{P}))$.*

For the remainder of this section let $\vec{U}$, $\vec{S}$, and $\mathcal{P}$ be as in Theorem 25.

We need the following terminology. A separation $\vec{s} \in \vec{S}$ is *exclusive (for $\mathcal{P}$)* if it lies in exactly one profile in $\mathcal{P}$. If $P \in \mathcal{P}$ is the profile 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 be 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 $\mathcal{P}$. These sets M_P may be empty. Our first lemma addresses this:

Lemma 4.1.10. *If $\mathcal{P}$ is non-empty, then some M_P is non-empty.*

The existence of exclusive separations is actually an immediate consequence of Theorem 16: if $N \subseteq S$ is a nested set which distinguishes $\mathcal{P}$, each element of which distinguishes some pair of profiles in $\mathcal{P}$, then each 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.

To avoid the proof of Theorem 25 relying on its non-canonical version, let us give an independent proof of Lemma 4.1.10.

Proof of Lemma 4.1.10. If $\mathcal{P}$ consists of only one profile the assertion is trivial. 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 profiles in $\mathcal{P}$ has two exclusive orientations, one in each profile.

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 profile and set $\mathcal{P}' := \mathcal{P} \smallsetminus \{P\}$. By the induction hypothesis applied to $\mathcal{P}'$ and an arbitrary profile 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 either of $\vec{r}$ and $\vec{s}$ 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 $\mathcal{P}$. Thus we may assume that r and s cross.

By submodularity of S one of $\vec{r} \vee \overleftarrow{s}$ and $\overleftarrow{r} \vee \vec{s}$ lies in $\vec{S}$; by symmetry we may assume that $(\vec{r} \vee \overleftarrow{s}) \in \vec{S}$. Since $\vec{s}$ is Q'-exclusive we have $\overleftarrow{s} \in Q$ and hence $(\vec{r} \vee \overleftarrow{s}) \in Q$ by the profile property. From $(\vec{r} \vee \overleftarrow{s}) \geqslant \vec{r}$ we infer that $(\vec{r} \vee \vec{s})$ 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 the stronger assertion used for the induction hypothesis in this proof, too, can be established immediately using Theorem 16: 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 25, let us show that separations from different M_P's cannot cross:

Lemma 4.1.11. *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 submodularity one of $\vec{r} \vee \overleftarrow{s}$ and $\overleftarrow{r} \vee \vec{s}$ lies 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 profile in $\mathcal{P}$ contains both $\overleftarrow{r}$ and $\overleftarrow{s}$ and would then be inconsistent.

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 profiles 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 4.1.12. *The separation system $\vec{S'}$ is submodular and distinguishes $\mathcal{P}'$.*

Proof. The fact that $\vec{S'}$ is submodular is immediate from Lemma 2.1.1. For the latter, let Q and Q' be distinct profiles 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 something 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 $\vec{r} \vee \vec{s'}$ cannot lie in $\vec{S}$ since that would be a strictly larger P-exclusive separation than $\vec{r}$. Hence $(\vec{r} \wedge \vec{s'}) \in \vec{S}$. By $P \notin \{Q, Q'\}$ we have that both Q and Q' contain $\overleftarrow{r}$, and hence this corner separation distinguishes Q and Q' as well. However, by Lemma 2.1.1 and Lemma 4.1.11, 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. $\square$

Let us now mend the fact that the sets M_P may not themselves be nested.

Lemma 4.1.13. *Assume that $|\mathcal{P}| \neq 1$. Then the infimum in $\vec{U}$ of each non-empty M_P is P-exclusive; in particular it lies in $P \subseteq \vec{S}$. Moreover, if some $t \in U$ is nested with M_P, then t is also nested with this infimum.*

Proof. Fix an enumeration $M_P = \{\vec{r}_1, \dots, \vec{r}_n\}$ and some $t \in U$ that is nested with M_P. For $i = 1, \dots, n$ let $\vec{s}_i := \vec{r}_1 \wedge \dots \wedge \vec{r}_i$. We show by induction on i that $\vec{s}_i$ is P-exclusive and nested with t; this yields the claim for $i = n$.

The case $i = 1$ is trivially true, so suppose that $i > 1$ and that $\vec{s}_{i-1} = (\vec{r}_1 \wedge \dots \wedge \vec{r}_{i-1})$ is already known to be P-exclusive and nested with t.

If $s_{i-1} = r_i$ there is nothing to show, 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 $\vec{s}_i = (\vec{r}_i \wedge \vec{s}_{i-1})$ equals one of the two and hence is 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 profile 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 $\vec{r}_i \vee \vec{s}_{i-1}$ cannot lie in $\vec{S}$ since it would be P-exclusive and strictly larger than $\vec{r}_i \in M_P$. Therefore $(\vec{r} \wedge \vec{s}_{i-1}) \in \vec{S}$, that is, $\vec{s}_i \in \vec{S}$. By consistency we have that $\vec{s}_i \in P$. Every profile other than P contains $\overleftarrow{r}_i$ as well as $\overleftarrow{s}_{i-1}$ and hence $\overleftarrow{s}_{i-1}$ by the profile property, which shows that $\vec{s}_i$ is P-exclusive. Finally, by Lemma 2.1.1, $\vec{s}_i$ is also nested with t. $\square$

If M_P is non-empty let us write $\vec{s}_P$ for its infimum in $\vec{U}$. We are now ready to prove Theorem 25.

Proof of Theorem 25. 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 $\mathcal{P}$. Let

$$N_1 := \{s_P \mid M_P \neq \emptyset\} \;;$$

this is clearly a canonical set. By Lemma 4.1.13 N_1 distinguishes all profiles in $\mathcal{P}' \setminus \mathcal{P}$ from each other and from each profile in $\mathcal{P}'$.

By Lemma 4.1.11 every element of M_P is nested with every element of $M_{P'}$ for all $P \neq P'$. Applying the 'moreover'-part of Lemma 4.1.13 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 Lemma 4.1.10 and 4.1.12, yielding a canonical nested set $N_2 \subseteq S'$ which distinguishes $\mathcal{P}'$. The union $N_1 \cup N_2$ is then the desired nested set. $\square$

4.1.8 A short proof using unravelling

Let us briefly demonstrate why the ability to unravel submodular separation systems as studied in Section 3.6 would be useful by giving a short proof of Theorem 16 for such separation systems. Let U be a finite universe, $S \subseteq U$ submodular, and $\mathcal{P}$ a set of profiles of S. Assuming that S has an unravelling $S = \{s_1, \dots, s_n\}$, let us show that the assertion of Theorem 16 holds, i.e. that there is a nested set $N \subseteq S$ which distinguishes $\mathcal{P}$. For this we use induction on $n = |S|$. The case of $n = 1$ is clear, so suppose that $n > 1$ and that the assertion holds for each smaller value of n.

Let $\mathcal{P}' = \{P \smallsetminus s_n \mid P \in \mathcal{P}\}$. If there are no two profiles P_1 and P_2 in $\mathcal{P}$ with $P_1 \Delta P_2 = s_n$ we are done after applying the induction hypothesis to $S - s_n$ and $\mathcal{P}'$. So suppose that there are profiles P_1 and P_2 in $\mathcal{P}$ which differ only on s_n. We claim that the set S' of all separations in $S - s_n$ that are nested with s_n is submodular and distinguishes $\mathcal{P}'$. If this is true then we are done after applying the induction hypothesis to S' and $\mathcal{P}'$ and then adding s_n to the resulting nested set. So let us prove this claim.

Observe first that $S - s_n$ is submodular by unravelling, and hence $S' \subseteq U$ is submodular by Lemma 2.1.1. Let Q_1' and Q_2' be two profiles in $\mathcal{P}'$, distinguished by some $r \in S - s_n$. Suppose that r crosses s_n. Some orientation $\vec{r}$ of r lies in both P_1 and P_2. Then neither of $\vec{r} \vee \vec{s}_n$ and $\vec{r} \vee \overleftarrow{s}_n$ lies in $\vec{S}$, since these would be separations other than s_n distinguishing P_1 and P_2 by property P. Hence $\vec{r} \wedge \vec{s}_n$ and $\vec{r} \wedge \overleftarrow{s}_n$ lie in $\vec{S}$. In fact, being nested with s_n, they lie in $\vec{S'}$. One of these two corner separations distinguishes Q_1' and Q_2' by property P, concluding the proof. $\square$

Of course, the above approach only gives a full proof of Theorem 16 if it is true that every submodular separation system has an unravelling. Moreover the choice of this unravelling is non-canonical, and there is therefore no hope to obtain a proof of Theorem 25 from the previous section in this fashion.

4.2 The tangle-tree duality theorem

The second pillar of tangle theory consists of just a single result, the *tangle-tree duality theorem*:

Theorem 1 (Tangle-tree duality theorem [19]). *Let U be a finite universe of separations, $S \subseteq U$ 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 $\mathcal{F}$-separable. Then precisely one of the following holds:*

- *there is an S-tree over $\mathcal{F}$;*
- *there is an $\mathcal{F}$-tangle of S.*

The great strength of Theorem 1 lies in the flexibility of the choice of $\mathcal{F}$. This set $\mathcal{F}$ can be tailored to capture a wide variety of tangles and clusters, allowing Theorem 1 to be employed in a multitude of different settings ([16, 20]). Moreover the freedom of choosing and manipulating $\mathcal{F}$ lends itself to utilizing Theorem 1 to prove results which, on the surface, have nothing to do with tangle-tree duality: by clever choice of $\mathcal{F}$ one can ensure that there is no $\mathcal{F}$-tangle of S, and that the S-tree over $\mathcal{F}$ one the obtains has certain desirable properties. We shall see one such 'unforeseen' application of Theorem 1 in Section 4.3.

The original proof of Theorem 1 from [19] proceeds by induction on the number of $s \in S$ for which neither $\{\vec{s}\}$ nor $\{\overleftarrow{s}\}$ lies in $\mathcal{F}$. This results in a proof that is reasonably compact but on the other hand somewhat technical and not the most enlightening. Moreover with this approach one has no control over the set of separations in S for which one might have to invoke the assumption that S is $\mathcal{F}$-separable.

In this section we present two new proofs of Theorem 1, both of which improving in some aspects over the original. The core argument of both proofs is somewhat similar: if there is no S-tree over $\mathcal{F}$, then every 'attempt' at such an S-tree must fail. Thus, if one starts with some star σ in $\mathcal{F}$ as the basis for such an S-tree, and then 'glues' for each $\vec{s} \in \sigma$ with $\{\overleftarrow{s}\} \notin \mathcal{F}$ some star from $\mathcal{F}$ onto $\vec{s}$ that contains $\overleftarrow{s}$, then one must at some point be unable to find such a star. The resulting attempt is then an S-tree that is 'over $\mathcal{F}$' only for all internal vertices, but not necessarily at the leaves. The fundamental strategy of both proofs presented here is to collect the set of all these leaf-separations at which the S-tree attempts get stuck, and then turn this set into the basis of an $\mathcal{F}$-tangle.

The strategy described here is already present in Mazoit's proof ([39]) of the classical duality theorem for brambles and tree width in graphs. In [10] Diestel gives a proof of this graph-theoretic duality theorem that is derived from his and Oum's original proof of Theorem 1, applied to the specific $\mathcal{F}$ corresponding to tree decompositions of a certain width. Curiously Mazoit's and Diestel's graph-theoretic proofs are quite similar. One could therefore argue that the strategy for the proofs here comes from re-translating Diestel's translation of Theorem 1 to this specific graph application back into abstract separation systems.

Both of the proofs given here allow one to strengthen Theorem 1: they both allow a weakening of the technical assumption that S be $\mathcal{F}$-separable to only require $\mathcal{F}$-separability for those separations whose inverse lies in no star of $\mathcal{F}$, rather than for all separations in $\vec{S}$. This, of course, yields a stronger and more widely applicable version of Theorem 1. In fact, this strengthening enables an application that was previously not possible: in Section 4.3 we will see that a clever use of the tangle-tree duality theorem allows one to derive tree-of-tangles theorems from it. We shall obtain a variant of Theorem 16 there. Going even further, in [24] the authors utilise the tangle-tree duality theorem to prove *efficient* tree-of-tangles theorems, and in particular show a non-canonical version of Theorem 20. This latter application requires the strengthened tangle-tree duality theorem.

The first proof of Theorem 1 shown is a straightforward implementation of the base argument outlined above. It utilizes the same technical preliminaries as the original proof in [19], including in particular Lemma 2.5.1. The proof of Theorem 1 itself is only slightly shorter than the original. The main upside over the original proof, apart from giving a stronger statement, is that the proof given here is cleaner and less technical: the proof from [19], by design, has to deal with a plethora of corner cases in which the separations at hand might be trivial or degenerate. These distractions vanish almost completely here, since the separations collected from the leaves of 'failed' S-trees will automatically be nontrivial and nondegenerate.

The second proof presented here aims to eliminate the need for those technical preliminary lemmas from [19] such as Lemma 2.5.1. The way to do this is to get around the need to 'shift' entire S-trees, a technique required in both the original and the above first proof. This is achieved by expanding the set $\mathcal{F}$ without altering the outcome of Theorem 1's dichotomy. A side effect of this approach is that it allows an even further weakening of the assumption of $\mathcal{F}$-separability compared to the first proof, albeit a much less significant one than the improvement achieved in the first proof over [19].

Let us re-iterate that for this section, as well as for its successor Section 4.3, we assume that the reader is familiar with [19] in its entirety. We will be using definitions, notation, and techniques from [19] without introducing them explicitly; the definitions and tools given in Section 2.5 cover only the basics. In particular we shall assume familiarity with the following from [19] without further explanation: *tight*, *irredundant*, and *order-respecting* S-trees; stars associated with nodes of S-trees; the *shifting map* $f\downarrow^{\vec{r}}_{\vec{s}_0}$ as well as all surrounding notation; *separable* and *closed under shifting*; and the shift $\alpha_{r|v_r,\vec{s}_0}$ of an S-tree. We will also use a series of lemmas from [19], which we will point out explicitly.

4.2.1 First proof: Petals

Let us dive straight into the first proof of Theorem 1. For the remainder of this section 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. We may assume that $\emptyset \notin \mathcal{F}$.

We shall need the following definitions. An *S-tree attempt* is an S-tree (T, α) with at least one edge and $\alpha(t) \in \mathcal{F}$ for every internal node t of T. For a leaf t of an S-tree attempt (T, α) the *incoming label* of t is the separation $\vec{r}$ for which $\{\vec{r}\}$ is associated with t in (T, α); the *outgoing label* of t is then $\overleftarrow{r}$. We call $\vec{r}$ a *petal* of (T, α) if $\{\vec{r}\} \notin \mathcal{F}$.

A separation $\vec{r}$ in $\vec{S}$ is *$\mathcal{F}$-critical* if $\vec{r} \in \sigma$ for some $\sigma \in \mathcal{F}$, but there is no $\sigma' \in \mathcal{F}$ with $\sigma' \cap r = \{\overleftarrow{r}\}$. Observe that if $\vec{r} \in \vec{S}$ is $\mathcal{F}$-critical then $\vec{r}$ is nondegenerate and not forced by $\mathcal{F}$, and in particular $\vec{r}$ is nontrivial in S

since $\mathcal{F}$ is standard for S. We say that S is *critically $\mathcal{F}$-separable* if for all $\mathcal{F}$-critical $\vec{r}, \overleftarrow{r'} \in \vec{S}$ with $\vec{r} \leqslant \vec{r'}$ there exists an $s_0 \in S$ with an orientation $\vec{s}_0$ that emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$ and such that $\overleftarrow{s}_0$ emulates $\overleftarrow{r'}$ in $\vec{S}$ for $\mathcal{F}$. Clearly, if S is $\mathcal{F}$-separable, then S is critically $\mathcal{F}$-separable.

Assume for the remainder of this section that S is critically $\mathcal{F}$-separable. We will prove the following assertion which is equivalent to the tangle-tree duality theorem:

Proposition 4.2.1. *S has an $\mathcal{F}$-tangle if and only if every S-tree attempt has a petal.*

Since the S-tree attempts without petals are exactly the S-trees that are over $\mathcal{F}$, Proposition 4.2.1 immediately implies Theorem 1.

For our proof of Proposition 4.2.1 we will use [19, Lemmas 2.1–2.4], or equivalently, their original sources [13, Lemma 6.2–6.5]. These lemmas, roughly speaking, say that an S-tree over a set of stars may be assumed to be 'cleaned up', i.e. tight and irredundant; that in such a cleaned up S-tree (T, α) the map α is a homomorphism; and that then each nontrivial separation can only appear once as a label. In short, cleaning up a given S-tree over stars enables us to apply Lemma 2.5.1 to it.

Proof of Proposition 4.2.1. The forward direction of Proposition 4.2.1 is clear, so let us show the backward direction.

Let $\mathcal{P} \subseteq \vec{S}$ be the set of all petals of S-tree attempts. Suppose first that we can find $P \subseteq \mathcal{P}$ such that P is a consistent and antisymmetric set that contains at least one petal of every S-tree attempt. Then P is an $\mathcal{F}$-tangle of S: P orients each separation s in S since it contains a petal of the S-tree attempt that is just a single edge labelled with s; and P avoids each star σ in $\mathcal{F}$ since it is antisymmetric and contains a petal of the S-tree attempt consisting of one internal node for σ and one leaf for every element of σ.

So let us show that we can find such a set $P \subseteq \mathcal{P}$. For this pick a $P \subseteq \mathcal{P}$ that is minimal with respect to inclusion subject to the conditions that P contains at least one petal of every S-tree attempt and is down-closed in $\mathcal{P}$, that is, such that $\vec{p} \in P$ for all $\vec{p} \in \mathcal{P}$ with $\vec{p} \leqslant \vec{q}$ for some $\vec{q} \in P$. Such a P exists since $\mathcal{P}$ itself is a candidate. We claim that this P is antisymmetric and consistent.

So suppose that P is not antisymmetric, or not consistent. Then there are $\vec{r} \neq \vec{s}$ in P with $\overleftarrow{r} \leqslant \vec{s}$. In particular we can take $\vec{r}$ and $\vec{s}$ to be maximal elements of P. Neither $\vec{r}$ nor $\vec{s}$ can be co-trivial in S since $\mathcal{F}$ is standard for S and $\vec{r}$ and $\vec{s}$ are petals. Therefore neither of the two can be trivial or degenerate either, since this would imply that the other one is co-trivial.

By picking $\vec{r}$ and $\vec{s}$ among the maximal elements of P we ensure that both $P \setminus \{\vec{r}\}$ and $P \setminus \{\vec{s}\}$ are still down-closed in $\mathcal{P}$. Thus, by the minimality of P, there are S-tree attempts (T_r, α_r) and (T_s, α_s) whose only petals that lie in P are $\vec{r}$ and $\vec{s}$, respectively. We may assume (T_r, α_r) and (T_s, α_s) to be tight and irredundant (cf. [13, Lemmas 6.2 and 6.5]), which implies that $\vec{r}$ and $\vec{s}$ are the incoming label of exactly one leaf of T_r and T_s, respectively (cf. [13, Lemma 6.4]).

We claim that $\overleftarrow{r}$ is $\mathcal{F}$-critical. To see this, let v_r be the leaf of T_r whose incoming edge is labelled with $\vec{r}$, and let w_r be the neighbour of v_r in T_r. Then w_r has an incoming edge labelled with $\overleftarrow{r}$, and we must have $\alpha_r(w_r) \in \mathcal{F}$, witnessing that $\overleftarrow{r}$ lies in some star in $\mathcal{F}$: for if $\alpha_r(w_r) \notin \mathcal{F}$ then w_r would be a leaf of T_r and $\overleftarrow{r}$ a petal of (T_r, α_r). By $\overleftarrow{r} \leqslant \vec{s}$ and P being down-closed in $\mathcal{P}$ we would then have $\overleftarrow{r} \in P$, contrary to the assumption that $\vec{r}$ is the only petal of (T_r, α_r) which lies in P. Suppose now that $\sigma \cap r = \{\vec{r}\}$ for some $\sigma \in \mathcal{F}$. Then we can extend (T_r, α_r) by σ at the leaf at which $\vec{r}$ appears. This extension of (T_r, α_r) is then an S-tree attempt, of which P must contain a petal. Since $\vec{r} \notin \sigma$ this petal would be strictly larger than $\vec{r}$, contradicting the maximality of $\vec{r}$ in P.

A similar argument shows that $\overleftarrow{s}$ is $\mathcal{F}$-critical. By the assumption that S is critically $\mathcal{F}$-separable we thus find an $s_0 \in S$ with an orientation $\vec{s}_0$ that emulates $\overleftarrow{r}$ in $\vec{S}$ for $\mathcal{F}$ and such that $\overleftarrow{s}_0$ emulates $\overleftarrow{s}$ in $\vec{S}$ for $\mathcal{F}$. Let v_r and v_s be the leaves of T_r and T_s with incoming labels $\vec{r}$ and $\vec{s}$, respectively. Set $\alpha'_r := \alpha_{r|v_r,\vec{s}_0}$ and $\alpha'_s := \alpha_{s|v_s,\overleftarrow{s}_0}$. Then Lemma 2.5.1 says that (T_r, α'_r) and (T_s, α'_s) are S-tree attempts in which v_r is the unique leaf of T_r with incoming label $\overleftarrow{s}_0$, and that v_s is the unique leaf of T_s with incoming label $\vec{s}_0$. Let (T, α) be the S-tree obtained from (T_r, α'_r) and (T_s, α'_s) by identifying $v_r \in T_r$ with the neighbour of v_s in T_s and vice-versa, and extending the maps α'_r and α'_s accordingly.

This (T, α), too, is an S-tree attempt. Let v_r be the leaf of (T_r, α_r) that is associated with $\vec{r}$, and likewise v_s the leaf of (T_s, α_s) associated with $\vec{s}$. Observe that every leaf of T is either a leaf of T_r other than v_r, or a leaf of T_s other than v_s.

We claim that P, contrary to its definition, contains no petal of (T, α). To see this, let t be some leaf of T_r other than v_r, and let $\vec{p}$ be the incoming label of t in (T_r, α_r). Then $\overleftarrow{r} \leqslant \vec{p}$ (cf. [13, Lemma 6.3]), and thus the incoming label of t in (T, α) is $\vec{s}_0 \vee \vec{p}$. If $\alpha_r(t) = \{\vec{p}\} \in \mathcal{F}$ then $\alpha(t) = \{\vec{s}_0 \vee \vec{p}\} \in \mathcal{F}$ since $\vec{s}_0$ emulates $\overleftarrow{r}$ in $\vec{S}$ for $\mathcal{F}$. Consequently, if $\vec{s}_0 \vee \vec{p}$ is a petal of (T, α), then $\vec{p}$ is a petal of (T_r, α_r). In particular, since $\vec{p} \neq \vec{r}$ and $\vec{r}$ is the only petal of P from (T_r, α_r), and P is down-closed in $\mathcal{P}$, we know that $(\vec{s}_0 \vee \vec{p}) \geqslant \vec{p}$ cannot be a petal of (T, α) that lies in P.

By the same argument those leaves of T which are leaves of T_s other than v_s cannot give rise to a petal of (T, α) in P, either. Hence P contains no petal from (T, α), causing a contradiction. This finishes the proof that P is consistent and antisymmetric and hence an $\mathcal{F}$-tangle of S. □

We have thus established the following strengthening of Theorem 1:

Theorem 26. *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.*

4.2.2 Second proof: Elimination

Our second proof of Theorem 1 still allows for a significant weakening of Theorem 1's assumptions, but mainly focuses on cutting down on the amount of technical preliminaries required to establish its assertion. In particular we shall be able to do without Lemma 2.5.1 and [13, Lemma 6.2–6.5], which we still needed in the previous section.

As in the previous section, 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. On a high level, the fundamental strategy of this second proof is the same as the first: we collect the set of all separations at leaves of S-tree attempts that are not, as singletons, included in $\mathcal{F}$, and then construct an $\mathcal{F}$-tangle from those separations. However, we will not handle these S-tree attempts directly. Instead we will alter the set $\mathcal{F}$ so that it includes, for each two stars in $\mathcal{F}$ that are 'adjacent' in the sense of containing different orientations of some separation, their 'elimination': the star obtained from their union by deleting the separation they orient differently. It will not be difficult to show that closing $\mathcal{F}$ under this operation influences neither whether S is $\mathcal{F}$-separable nor whether an S-tree over $\mathcal{F}$ exists. The effect of enriching $\mathcal{F}$ in this way is that for every S-tree attempt (T, α) the set of outgoing labels of its leaves will be a star in $\mathcal{F}$. For the final part of the proof we will therefore no longer need to call upon Lemma 2.5.1, instead shifting stars one at a time, and neither shall we need to glue together two S-trees, since this will also be handled by the 'elimination'-operation.

Let us make the above ideas formal. We say that a separation $\vec{s}_0 \in \vec{S}$ *symmetrically emulates* $\vec{r} \in \vec{S}$ *for* $\mathcal{F}$ if $\vec{s}_0$ emulates $\vec{r}$ in $\vec{S}$ and for every antisymmetric star $\sigma \subseteq \vec{S}_{\geqslant \vec{r}} \smallsetminus \{\overleftarrow{r}\}$ in $\mathcal{F}$ that has an element $\vec{r'} \geqslant \vec{r}$ the star $f \downarrow^{\vec{r}}_{\vec{s}_0} (\sigma)$ either lies in $\mathcal{F}$ or fails to be antisymmetric. (This definition is exactly the same as the usual 'emulates for $\mathcal{F}$', except that we only ask that $f \downarrow^{\vec{r}}_{\vec{s}_0} (\sigma) \in \mathcal{F}$ for stars where both σ and its image are antisymmetric.) Recall that $\vec{r}$ in $\vec{S}$ is $\mathcal{F}$*-critical* if $\vec{r} \in \sigma$ for some $\sigma \in \mathcal{F}$, but there is no $\sigma' \in \mathcal{F}$ with $\sigma' \cap r = \{\overleftarrow{r}\}$.

We say that S is *symmetrically* $\mathcal{F}$*-separable* if for all $\mathcal{F}$-critical $\vec{r}, \overleftarrow{r'} \in \vec{S}$ with $\vec{r} \leqslant \vec{r'}$ there exists an $s_0 \in S$ with an orientation $\vec{s}_0$ that symmetrically emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}$ and such that $\overleftarrow{s}_0$ symmetrically emulates $\overleftarrow{r'}$ in $\vec{S}$ for $\mathcal{F}$. Since $\mathcal{F}$ is standard for S these $\mathcal{F}$-critical separations can be neither trivial nor degenerate. Clearly, if S is $\mathcal{F}$-separable then S is symmetrically $\mathcal{F}$-separable.

The strengthening of Theorem 1 we are establishing then reads as follows:

Theorem 27. *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 symmetrically $\mathcal{F}$-separable. Then precisely one of the following holds:*

- *there is an S-tree over $\mathcal{F}$;*
- *there is an $\mathcal{F}$-tangle of S.*

Let us henceforth assume that S is symmetrically $\mathcal{F}$-separable. For our proof of Theorem 27 we will be relying on the following two concepts.

Given stars σ and ρ, we call σ an *expansion* of ρ if $\sigma \supseteq \rho$ with σ being antisymmetric and $\sigma \smallsetminus \rho$ consisting solely of separations that are trivial in S. If s is a separation in S and $\sigma_{\vec{s}}, \sigma_{\overleftarrow{s}}$ are stars with $\vec{s} \in \sigma_{\vec{s}}$ and $\overleftarrow{s} \in \sigma_{\overleftarrow{s}}$, the star $\sigma := (\sigma_{\vec{s}} \smallsetminus \{\vec{s}\}) \cup (\sigma_{\overleftarrow{s}} \smallsetminus \{\overleftarrow{s}\})$ is called an *elimination* of $\sigma_{\vec{s}}$ and $\sigma_{\overleftarrow{s}}$.

The first part of our proof of Theorem 27 consists of showing that we may assume, essentially, that $\mathcal{F}$ consists of antisymmetric stars and is closed under expansion and elimination. For this observe first that expansions of stars are antisymmetric by definition, and that the elimination of two antisymmetric stars $\sigma_{\vec{s}}$ and $\sigma_{\overleftarrow{s}}$ is again antisymmetric: for if there were a separation $t \neq s$ with, say, $\vec{t} \in \sigma_{\vec{s}}$ and $\overleftarrow{t} \in \sigma_{\overleftarrow{s}}$, we would have $\vec{t} \leqslant \overleftarrow{s}$ as well as $\overleftarrow{s} \leqslant \vec{t}$ by the star property in those respective stars, giving $\vec{t} = \overleftarrow{s}$ in violation of $\sigma_{\vec{s}}$'s antisymmetry. In fact, if $\sigma_{\vec{s}}$ and $\sigma_{\overleftarrow{s}}$ are antisymmetric, their elimination σ equals $(\sigma_{\vec{s}} \cup \sigma_{\overleftarrow{s}}) \smallsetminus \{\vec{s}, \overleftarrow{s}\}$, and every nontrivial element of σ lies in exactly one of $\sigma_{\vec{s}}$ or $\sigma_{\overleftarrow{s}}$.

Let $\mathcal{F}_0 \subseteq \mathcal{F}$ consist of all antisymmetric stars in $\mathcal{F}$. For $n \geqslant 1$ let $\mathcal{F}_n$ be the set of all stars σ in S that are an expansion or an elimination of star(s) in $\mathcal{F}_{n-1}$. Note that $\mathcal{F}_{n-1} \subseteq \mathcal{F}_n$ since each $\sigma \in \mathcal{F}_{n-1}$ is an expansion of itself. Further let $\mathcal{F}' := \bigcup_{n\in\mathbb{N}} \mathcal{F}_n$; this is a set of antisymmetric stars that is closed under expansion and elimination. Clearly $\mathcal{F}'$ is standard for S.

Our first lemma says that replacing $\mathcal{F}$ with $\mathcal{F}'$ does not change whether there is an S-tree over that set of stars:

Lemma 4.2.2. *There is an S-tree over $\mathcal{F}$ if and only if there is an S-tree over $\mathcal{F}'$.*

Proof. For the forward direction observe that any S-tree that is over $\mathcal{F}$ and, subject to that, contains as few nodes as possible is over $\mathcal{F}'$: any node whose associated star fails to be antisymmetric can be suppressed. Alternatively one may use [13, Lemmas 6.5], which shows slightly more.

For the backward direction, since $\mathcal{F}' = \mathcal{F}_n$ for some $n \in \mathbb{N}$, it suffices to show that there is an S-tree over $\mathcal{F}_{n-1}$ whenever there is an S-tree over $\mathcal{F}_n$, for $n \geqslant 1$. So suppose that (T, α) is an S-tree over $\mathcal{F}_n$ with as few nodes as possible whose associated star does not lie in $\mathcal{F}_{n-1}$. We claim that (T, α) is over $\mathcal{F}_{n-1}$. Suppose not; then there is a node t of T whose associated star $\sigma = \alpha(t)$ lies in $\mathcal{F}_n \smallsetminus \mathcal{F}_{n-1}$, i.e. is an expansion or elimination of some star(s) in $\mathcal{F}_{n-1}$. Let us treat these two cases in turn.

If σ is an expansion of some $\rho \in \mathcal{F}_{n-1}$, we can delete from T all components of $T - t$ whose outgoing edge to t is labelled with an element of $\sigma \smallsetminus \rho$, thereby obtaining an S-tree which has (at least) one fewer node than (T, α) whose associated star fails to lie in $\mathcal{F}_{n-1}$.

On the other hand, if σ is the elimination of some $\sigma_{\vec{s}}, \sigma_{\overleftarrow{s}} \in \mathcal{F}_{n-1}$, we can split the vertex t into two new vertices $t_{\vec{s}}$ and $t_{\overleftarrow{s}}$ joined by an edge labelled with s such that their associated stars are $\sigma_{\vec{s}}$ and $\sigma_{\overleftarrow{s}}$, respectively. For this let t' be a neighbour of t in T and C the component of $T - t$ containing t'. Then $\vec{r} := \alpha(t', t)$ lies in $\sigma_{\vec{s}}$ or $\sigma_{\overleftarrow{s}}$. If $\vec{r}$ lies in exactly one of the two, say, in $\sigma_{\vec{s}}$, we let t' be adjacent with the corresponding vertex $t_{\vec{s}}$. If $\vec{r}$ lies in both of $\sigma_{\vec{s}}$ and $\sigma_{\overleftarrow{s}}$ then $\vec{r}$ is trivial with witness s. In this case we delete C from T and create two new vertices t_1 and t_2, make t_1 adjacent to $t_{\vec{s}}$, make t_2 adjacent to $t_{\overleftarrow{s}}$, and label the outgoing edge of both t_1 and t_2 with $\vec{r}$. Since $\mathcal{F}_{n-1}$ is standard for S and thus contains $\{\overleftarrow{r}\}$ for all such $\vec{r}$, the resulting S-tree will have at least one fewer vertex than (T, α) whose associated star does not lie in $\mathcal{F}_{n-1}$. $\square$

The next step is to show that S is symmetrically $\mathcal{F}'$-separable.

Lemma 4.2.3. *If S is symmetrically $\mathcal{F}$-separable then it is symmetrically $\mathcal{F}'$-separable.*

Proof. We show that S is symmetrically $\mathcal{F}_0$-separable, and then for $n \geqslant 1$ that S is symmetrically $\mathcal{F}_n$-separable whenever it is symmetrically $\mathcal{F}_{n-1}$-separable.

For the first let $\vec{r}, \overleftarrow{r'} \in \vec{S}$ be $\mathcal{F}_0$-critical with $\vec{r} \leqslant \overrightarrow{r'}$. Then both $\vec{r}$ and $\overleftarrow{r'}$ are nontrivial, and consequently neither of them is co-trivial either, since that would imply that the other one is trivial. It suffices to check that $\vec{r}$ and $\overleftarrow{r'}$ are also $\mathcal{F}$-critical; we shall do this for $\vec{r}$ only. We need to show that there is no $\sigma \in \mathcal{F}$ with $\sigma \cap r = \{\overleftarrow{r}\}$. So suppose there is such a star $\sigma \in \mathcal{F}$. Then $\sigma \notin \mathcal{F}_0$ since $\vec{r}$ is $\mathcal{F}_0$-critical. Thus by definition of $\mathcal{F}_0$ we know that σ is not antisymmetric, i.e. contains $\vec{s}$ and $\overleftarrow{s}$ for some $s \in S$. By $\sigma \cap r = \{\overleftarrow{r}\}$ we must have $r \neq s$. But then $\overleftarrow{r}$ is trivial, a contradiction.

Suppose now that S is symmetrically $\mathcal{F}_{n-1}$-separable for $n \geqslant 1$ and let us show that is is symmetrically $\mathcal{F}_n$-separable. For this let $\vec{r}, \overleftarrow{r'} \in \vec{S}$ be $\mathcal{F}_n$-critical with $\vec{r} \leqslant \overrightarrow{r'}$. We must verify that $\vec{r}$ and $\overleftarrow{r'}$ are also $\mathcal{F}_{n-1}$-critical. We shall do this for $\vec{r}$; the proof for $\overleftarrow{r'}$ is similar. Since $\vec{r}$ is $\mathcal{F}_n$-critical and $\mathcal{F}_{n-1} \subseteq \mathcal{F}_n$ there is no $\sigma \in \mathcal{F}_{n-1}$ with $\sigma \cap r = \{\overleftarrow{r}\}$. To see that $\vec{r}$ lies in some star in $\mathcal{F}_{n-1}$ let $\sigma \in \mathcal{F}_n$ be a star with $\vec{r} \in \sigma$. If σ is an expansion of some $\rho \in \mathcal{F}_{n-1}$ then $\vec{r} \in \rho$ since $\vec{r}$ is nontrivial. Otherwise σ is the elimination of some $\sigma_{\vec{s}}$ and $\sigma_{\overleftarrow{s}}$ in $\mathcal{F}_{n-1}$. But then $\vec{r}$ is contained in one of the two and we are done.

We have shown that $\vec{r}$ and $\overleftarrow{r'}$ are $\mathcal{F}_{n-1}$-critical and can hence use that S is symmetrically $\mathcal{F}_{n-1}$-separable. Let $s_0 \in S$ be a separation with an orientation $\vec{s}_0$ that symmetrically emulates $\vec{r}$ in $\vec{S}$ for $\mathcal{F}_{n-1}$ and such that $\overleftarrow{s}_0$ symmetrically emulates $\overleftarrow{r'}$ in $\vec{S}$ for $\mathcal{F}_{n-1}$. We show that $\vec{s}_0$ and $\overleftarrow{s}_0$ do so for $\mathcal{F}_n$ as well. Let us show this for $\vec{s}_0$; the proof for $\overleftarrow{s}_0$ is analogous.

For readability we write $f := f\downarrow^{\vec{r}}_{\vec{s}_0}$. Let $\sigma \subseteq \vec{S}_{\geqslant \vec{r}} \smallsetminus \{\overleftarrow{r}\}$ be a star in $\mathcal{F}_n$ and let $\vec{t}$ be the unique element of σ with $\vec{r} \leqslant \vec{t}$. If $f(\sigma)$ is not antisymmetric there is nothing to show, so let us assume that it is. We need the show that $f(\sigma)$ lies in $\mathcal{F}_n$.

Consider first the case that σ is an expansion of some $\rho \in \mathcal{F}_{n-1}$. Then $\vec{t} \in \rho$ since $\vec{r}$ and hence $\vec{t}$ are nontrivial. Moreover $\rho \subseteq \sigma \subseteq \vec{S}_{\geqslant \vec{r}} \smallsetminus \{\overleftarrow{r}\}$, and hence $f(\rho) \in \mathcal{F}_{n-1}$ by choice of $\vec{s}_0$ since $f(\rho) \subseteq f(\sigma)$ is antisymmetric. For each $\vec{u} \in \sigma$ with $f(\vec{u}) \in f(\sigma) \smallsetminus f(\rho)$ we know that $\vec{u} \in (\sigma \smallsetminus \rho)$ is trivial in S, giving $t \neq r$ and thus $f(\vec{t}) \leqslant \vec{t}$ by definition of f. Therefore $f(\vec{t})$ is also trivial in S, showing that $f(\sigma)$ is an expansion of $f(\rho)$.

Let us now consider the case that σ is the elimination of $\sigma_{\vec{s}}, \sigma_{\overleftarrow{s}} \in \mathcal{F}_{n-1}$. By symmetry we may assume that $\vec{t} \in (\sigma_{\overleftarrow{s}} \smallsetminus \{\overleftarrow{s}\})$, which implies that $\vec{s}$ is the unique element of $\sigma_{\vec{s}}$ with $\vec{r} \leqslant \vec{s}$. Then both $\sigma_{\vec{s}}$ and $\sigma_{\overleftarrow{s}}$ are subsets of $\vec{S}_{\geqslant \vec{r}} \smallsetminus \{\overleftarrow{r}\}$: otherwise $r = s$, which would imply either $\vec{r} \leqslant \vec{t} \leqslant \vec{r}$ or $\vec{r} \leqslant \vec{t} \leqslant \overleftarrow{r}$, with the first violating the antisymmetry of $\sigma_{\overleftarrow{s}}$ and the latter the nontriviality of $\vec{r}$. Therefore $f(\sigma_{\vec{s}})$ and $f(\sigma_{\overleftarrow{s}})$ are well-defined and stars.

If both of $f(\sigma_{\vec{s}})$ and $f(\sigma_{\overleftarrow{s}})$ are antisymmetric then they lie in $\mathcal{F}_{n-1}$ by assumption, in which case $f(\sigma)$ is their elimination and hence lies in $\mathcal{F}_n$ as required.

So suppose that $f(\sigma_{\vec{s}})$ is not antisymmetric. Since $f(\sigma_{\vec{s}} \smallsetminus \{\vec{s}\}) \subseteq f(\sigma)$ is antisymmetric this means that $f(\overleftarrow{s})$ lies in $f(\sigma_{\vec{s}} \smallsetminus \{\vec{s}\}) \subseteq f(\sigma)$. Consequently every element of $f(\sigma_{\vec{s}})$ apart from $f(\vec{s})$ and $f(\overleftarrow{s})$ is trivial with witness $f(s)$. From $f(\overleftarrow{s}) \in f(\sigma)$ we infer that $f(\sigma_{\overleftarrow{s}}) \subseteq f(\sigma)$, giving $f(\sigma_{\overleftarrow{s}}) \in \mathcal{F}_{n-1}$ by the antisymmetry of $f(\sigma)$. Since every $\vec{u}$ in $f(\sigma) \smallsetminus f(\sigma_{\overleftarrow{s}})$ lies in $f(\sigma_{\vec{s}})$ with $u \neq f(s)$, and is hence trivial, we have that $f(\sigma)$ is an expansion of $f(\sigma_{\overleftarrow{s}}) \in \mathcal{F}_{n-1}$ and thus lies in $\mathcal{F}_n$.

Analogously, if $f(\sigma_{\overleftarrow{s}})$ is not antisymmetric, then $f(\sigma)$ is an expansion of $f(\sigma_{\vec{s}}) \in \mathcal{F}_{n-1}$. ☐

We remark that the proof of Lemma 4.2.3 becomes much shorter if one only wants to prove Theorem 1 and does not care for the differences between '$\mathcal{F}$-separable' and 'symmetrically $\mathcal{F}$-separable'.

We are now ready to prove Theorem 27.

Proof of Theorem 27. It is easy to see there there cannot be both an $\mathcal{F}$-tangle of S and an S-tree over $\mathcal{F}$. We must therefore show that at least one of the two exists. So suppose that there is no S-tree over $\mathcal{F}$. Then, by Lemma 4.2.2, there is no S-tree over $\mathcal{F}'$ either. We shall construct an $\mathcal{F}'$-tangle of S. This will then be an $\mathcal{F}$-tangle as well since it avoids all non-antisymmetric stars by definition.

Let $\mathcal{P}$ be the set of all $\vec{r} \in \vec{S}$ with $\{\vec{r}\} \notin \mathcal{F}'$. Suppose first that we can find $P \subseteq \mathcal{P}$ such that P is a consistent and antisymmetric set that contains, for each star in $\mathcal{F}'$, the inverse of some element of that star. Then P contains no co-trivial separation since $\mathcal{F}'$ is standard for S. We can thus apply Lemma 2.4.1 to P to obtain a consistent orientation $O \supseteq P$ of S. This O is then an $\mathcal{F}'$-tangle: for each star $\sigma \in \mathcal{F}'$ the orientation O contains some $\overleftarrow{s}$ with $\vec{s} \in \sigma$, and hence avoids σ.

So let us show that we can find such a set $P \subseteq \mathcal{P}$. For this pick a $P \subseteq \mathcal{P}$ that is minimal with respect to inclusion subject to the conditions that P contains for each star in $\mathcal{F}'$ the inverse of some element of that star, and that P is down-closed in $\mathcal{P}$, that is, such that $\vec{p} \in P$ for all $\vec{p} \in \mathcal{P}$ with $\vec{p} \leqslant \vec{q}$ for some $\vec{q} \in P$. Such a P exists since $\mathcal{P}$ itself is a candidate: if, for some $\sigma \in \mathcal{F}'$ there would be no $\overleftarrow{s} \in \mathcal{P}$ with $\vec{s} \in \sigma$, then the S-tree obtained from σ would be over $\mathcal{F}'$.

We claim that this P is antisymmetric and consistent. Suppose for a contradiction that this is not the case. Then there are $\vec{r} \neq \vec{s}$ in P with $\overleftarrow{r} \leqslant \vec{s}$. In particular we can take $\vec{r}$ and $\vec{s}$ to be maximal elements of P. Neither $\vec{r}$ nor $\vec{s}$ can be co-trivial in S since $\mathcal{F}'$ is standard for S and $\vec{r}, \vec{s} \in \mathcal{P}$. Therefore neither of the two can be trivial or degenerate either, since this would imply that the other one is co-trivial.

By picking $\vec{r}$ and $\vec{s}$ among the maximal elements of P we ensure that both $P \setminus \{\vec{r}\}$ and $P \setminus \{\vec{s}\}$ are still down-closed in $\mathcal{P}$. Thus, by the minimality of P, there are stars σ_r and σ_s in $\mathcal{F}'$ whose only elements with inverse in P are $\overleftarrow{r}$ and $\overleftarrow{s}$, respectively. By choosing σ_r and σ_s inclusion-minimal with this property we can ensure that $\{\overleftarrow{t}\} \notin \mathcal{F}'$ for each element $\vec{t}$ of those stars other than $\overleftarrow{r}$ and $\overleftarrow{s}$, respectively: for if $\vec{t} \in \sigma_r$ with $\vec{t} \neq \overleftarrow{r}$, say, and $\{\overleftarrow{t}\} \in \mathcal{F}'$, then the elimination $\sigma_r \setminus \{\vec{t}\}$ of σ_r and $\{\overleftarrow{t}\}$ still has $\overleftarrow{r}$ as the only element with $\vec{r} \in P$.

We claim that $\overleftarrow{r}$ is $\mathcal{F}$-critical. Since we have $\overleftarrow{r} \in \sigma_r$ it suffices to show that there is no $\sigma \in \mathcal{F}$ with $\sigma \cap r = \{\vec{r}\}$. Suppose that there is such a star σ. Consider the elimination $\rho := (\sigma_r \cup \sigma) \setminus r$ of these two stars. We have $\rho \in \mathcal{F}'$ by definition, and hence P contains some $\vec{t}$ with $\overleftarrow{t} \in \rho$. Then $t \neq r$, and by choice of σ_r we have $\overleftarrow{t} \notin \sigma_r$. But then $\overleftarrow{t} \in \sigma$, which also contains $\vec{r}$, giving $\vec{r} \leqslant \vec{t}$ by the star property. This contradicts the maximality of $\vec{r}$ in P.

A similar argument shows that $\overleftarrow{s}$ is $\mathcal{F}'$-critical. By the assumption that S is symmetrically $\mathcal{F}'$-separable we thus find an $s_0 \in S$ with an orientation $\vec{s}_0$ that symmetrically emulates $\overleftarrow{r}$ in $\vec{S}$ for $\mathcal{F}'$ and such that $\overleftarrow{s}_0$ symmetrically emulates $\overleftarrow{s}$ in $\vec{S}$ for $\mathcal{F}'$. Let $\sigma'_r := f \downarrow^{\overleftarrow{r}}_{\vec{s}_0} (\sigma_r)$ and $\sigma'_s := f \downarrow^{\overleftarrow{s}}_{\overleftarrow{s}_0} (\sigma_s)$. Since $\overleftarrow{r} \leqslant \vec{s}_0 \leqslant \vec{s}$ we have $\vec{s}_0 \in \sigma'_r$ and $\overleftarrow{s}_0 \in \sigma'_s$.

Suppose now that σ'_r is not antisymmetric. Since $\overleftarrow{r}$ is nontrivial and hence so is $\vec{s}_0 \geqslant \overleftarrow{r}$, we must have $s_0 \subseteq \sigma'_r$. There is thus some $\vec{t} \in \sigma_r$ with $\vec{t} \wedge \overleftarrow{s}_0 = f \downarrow^{\overleftarrow{r}}_{\vec{s}_0} (\vec{t}) = \overleftarrow{s}_0$. Then $\overleftarrow{t} \leqslant \vec{s}_0 \leqslant \vec{s}$, with the latter lying in P. By the minimality of σ_r we have $\overleftarrow{t} \in \mathcal{P}$. Therefore $\overleftarrow{t} \in P$ since P is down-closed in $\mathcal{P}$, contrary to the assumption that $\vec{r}$ is the only element of P whose inverse lies in σ_r.

A similar argument shows that σ'_s is antisymmetric. We hence have $\sigma'_r, \sigma'_s \in \mathcal{F}'$.

Let $\vec{t}$ be an element of σ_r other than $\overleftarrow{r}$. Then $\overleftarrow{t} \in \mathcal{P} \setminus P$ by minimality of σ_r. Since P is down-closed in $\mathcal{P}$, and $\overleftarrow{t} \leqslant \overleftarrow{t} \vee \vec{s}_0 = f \downarrow^{\overleftarrow{r}}_{\vec{s}_0} (\overleftarrow{t})$, the latter cannot lie in P either. Therefore P contains no separation whose inverse lies in $\sigma'_r \setminus s_0$.

A similar argument shows that P contains no separation whose inverse lies in $\sigma'_s \setminus s_0$ either. But then the elimination $\sigma' := (\sigma'_r \cup \sigma'_s) \setminus s_0$ is a star in $\mathcal{F}'$ with no element whose inverse lies in P, a contradiction. This concludes the proof that P is consistent and antisymmetric. $\square$

4.3 Merging the two pillars

In this section we show that through clever choice of $\mathcal{F}$ one can utilize Theorem 1 to prove a tree-of-tangles theorem for profiles in submodular separation systems, thereby demonstrating that the two pillars of tangle theory are not as independent from each other as they might at first appear.

We will go through the ideas of this type of proof in Section 4.3.1 and establish a basic tree-of-tangles theorem. A refined version of the same argument will be given in Section 4.3.2, where we show that the approach using tangle-tree duality yields some interesting benefits when compared to the proofs in the style of Section 4.1.

4.3.1 Trees-of-tangles from tangle-tree duality

The tree-of-tangles theorem we will prove using tangle-tree duality is the following:

Theorem 28. *Let $\vec{S}$ be a submodular separation system. Then S contains a nested set that distinguishes the set of regular profiles of S.*

By itself Theorem 28 is nothing special; indeed, it is a slight weakening of Theorem 16, which asserts the same but without requiring the profiles to be regular. In this case the proof of the pudding is not in the eating, but in its ingredients: we shall obtain Theorem 28 as a direct consequence of Theorem 1.

So let $\vec{S}$ be a 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 such that every element of $\mathcal{F}$ is included in at most one regular profile of S. If we can achieve this then Theorem 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 the set m_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 $m_P \in \mathcal{M} \subseteq \mathcal{F}$. Thus, if we could feed this $\mathcal{F}$ to Theorem 1, we would receive an S-tree over $\mathcal{F}$. The edge labels of this S-tree would be our desired nested set, since each element of $\mathcal{F}$ in included in at most one regular profile of S: indeed, the regular profiles of S have no subsets in $\mathcal{P}$ or $\mathcal{C}$, and each element $m_P \in \mathcal{M}$ is included only in P itself.

Unfortunately, we are still some way off from plugging $\mathcal{F}$ into Theorem 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 the following known result by Weißauer:

Lemma 4.3.1 ([16, Lemma 13])**.** *Let $\vec{U}$ be a universe of separations and $\vec{S} \subseteq \vec{U}$ a submodular separation system. Then $\vec{S}$ is separable.*

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 [14] a contingency plan was laid out for this exact situation: a series of lemmas from [14] 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 $\vec{S}$. 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 4.3.2 ([14, 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.

For a star $\sigma = \{\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n\} \subseteq \vec{S}$ a *shift of σ (to some $\vec{s}_0 \in \vec{S}$)* is a star $\sigma^* \subseteq \vec{S}$ of the form

$$\sigma^* = \{\vec{x}_1 \vee \vec{s}_0\,, \vec{x} \wedge \overleftarrow{s}_0, \dots, \vec{x}_n \wedge \overleftarrow{s}_0\}\ .$$

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 4.3.3 ([14, Lemma 13])**.** *If a regular profile of S includes a shift of some star, it also includes that star.*

In [14] the definition of shift of a star contains additional technical assumptions on σ and $\vec{s}_0$, keeping in line with the precise assumptions of Theorem 1. However the proof of Lemma 4.3.3 does not necessitate this, and neither does its application.

Lemma 4.3.3 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 star in $\mathcal{G}_{i-1}$. Write $\hat{\mathcal{F}}^* := \bigcup_{i \in \mathbb{N}} \mathcal{G}_i$. Then by Lemma 4.3.3 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. Lemmas 4.3.2 and 4.3.3 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 m_P, this implies that no other regular profile of S includes $\hat{\sigma}^*$.

We can thus formally prove Theorem 28:

Proof of Theorem 28. 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 $\mathcal{C} \subseteq \hat{\mathcal{F}}^*$. By Lemma 4.3.1 S is separable, and hence $\hat{\mathcal{F}}^*$-separable by construction of $\hat{\mathcal{F}}^*$. Hence we can apply the tangle-tree duality theorem 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 Lemmas 4.3.2 and 4.3.3 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 4.3.2, 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 1, which we may assume to be irredundant ([19, Lemma 2.3]). Let $\vec{N}$ be the image of α. Then N is a nested subset of S ([[19, Lemma 2.1]). 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 m_P, this $\hat{\sigma^*}$ must originate from m_P. Consequently no other regular profile of S can include $\hat{\sigma^*}$, since none of them include m_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 28. First, in the definition of $\mathcal{F}$, we could have used other sets $\mathcal{M}$: the only properties of $\mathcal{M}$ that we need are that no regular profile of S avoids $\mathcal{M}$, 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 the upcoming section, 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 28 to the level of Theorem 16 by dropping the assumption of regularity, since Lemma 4.3.3 cannot do without this regularity.

Finally, Elbracht and Teegen ([24]) have given a viable cart-before-the-horse version of the proof presented here: 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 those 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 1 one would only need to show two things, namely that $\mathcal{F}$ is closed under shifting and that there is no $\mathcal{F}$-tangle of S. The first of these amounts to Lemma 4.3.3; 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 4.3.2.

4.3.2 Degrees in trees-of-tangles

Let S be a structurally submodular separation system and P a regular profile of S. In this section we seek to answer the following question: over all trees-of-tangles that distinguish all regular profiles of S, how low can the degree of P in those trees-of-tangles be?

First, let us make this notion of degree in a tree-of-tangles formal. For this section 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 29. *Let S be a 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 29 we will follow the proof of Theorem 28, making a more refined choice of $\mathcal{M}$, and utilize the fact that uncrossing and shifting a set cannot increase its size.

We will later see an example of a submodular separation system in which $\delta(P) \leqslant 2$ for every profile P but $\Delta(T) = 3$ for every tree-of-tangles T; this will demonstrate that the last assertion of Theorem 29 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 29 is at most $|\max P|$.

Let us now prove Theorem 29:

Proof of Theorem 29. 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 28, 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 28 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, t the node of P in T, and $\hat{\sigma^*} := \alpha(t)$. As in the proof of Theorem 28 the only element of $\mathcal{F}$ from which $\hat{\sigma^*}$ can originate is d_P. Since uncrossing and shifting d_P cannot increase its size we have $|\hat{\sigma^*}| \leqslant |d_P| = \delta(P)$. Conversely we have $|\hat{\sigma^*}| \geqslant \delta(P)$ since $\hat{\sigma^*}$ distinguishes P from all other regular profiles. Thus the degree of P in (T, α) is indeed $\delta(P)$.

Finally, if $\Delta(T) > 3$, the maximum degree of T is attained in some node t whose associated star $\alpha(t)$ originates from some $d_P \in \mathcal{D}$, since all elements of $\hat{\mathcal{F}}^*$

originating from elements of $\mathcal{P}$ or $\mathcal{C}$ have size at most three. As above we thus have $|\alpha(t)| \leqslant |d_P| = \delta(P)$, giving $\Delta(T) = \delta(P)$. $\square$

Let us see an example showing that we cannot guarantee to find T with maximum degree less than three even if all regular profiles of S have $\delta(P) \leqslant 2$:

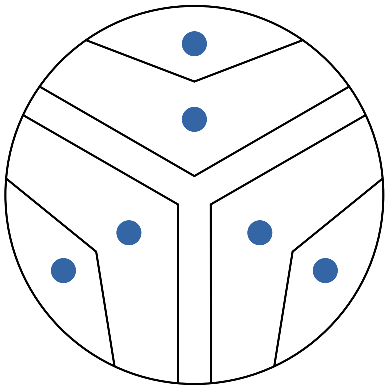

Figure 4.3: A ground-set and system of bipartitions.

Example 4.3.4. Let V consist of the six points in Fig. 4.3, and S be the separation system given by the six outlined bipartitions of V together with $\{\emptyset, V\}$. Then S is submodular and 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 maximum degree three.

Let us compare the tree-of-tangles found by Theorem 29 to the one outputted by the canonical tree-of-tangles theorem Theorem 25. For this let $\vec{S}$ be a submodular separation system and $\mathcal{P}$ the set of regular profiles of S. Evidently, for a regular profile $P \in \mathcal{P}$, there exists a P-exclusive separation if and only if $\delta(P) = 1$: the existence of such a P-exclusive separation witnesses that $\delta(P) \leqslant 1$, and conversely, if $\delta(P) = 1$, then the tree-of-tangles found by Theorem 29 contains a P-exclusive separation at some leaf. It is therefore the case that the tree-of-tangles constructed in Theorem 25 is optimal in terms of degrees for all profiles in $\mathcal{P}$ with $\delta(P) = 1$. However, for profiles $P \in \mathcal{P}$ with $\delta(P) > 1$ it is not necessarily the case that the tree-of-tangles from Theorem 25 optimizes their degree:

Example 4.3.5. Let V consist of the five points in Fig. 4.4, and S be the separation system given by the six outlined bipartitions of V together with $\{\emptyset, V\}$. Then S is submodular and the regular profiles of S correspond precisely to the five 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.

The tree-of-tangles found by Theorem 29 for this separation system consists of all regular separations in S, that is, of precisely the six bipartitions in Fig. 4.4. Let P be the profile induced by the centre point. Then $\delta(P) = 2$, and this is also the degree of P in this tree-of-tangles.

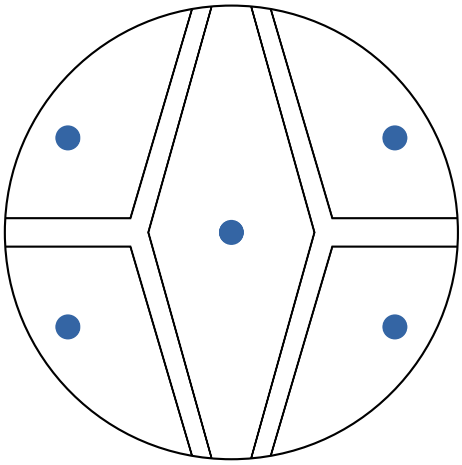

Figure 4.4: A ground-set and system of bipartitions.

In contrast to this, the nested set $N(\mathcal{P})$ given by Theorem 25 consists of only the four outer separations in Fig. 4.4: those separations in S that have an orientation exclusive to some profile in $\mathcal{P}$. The degree of P in the tree-of-tangles given by this set $N(\mathcal{P})$ is four and hence not optimal.

4.4 Tangles in graphs

4.4.1 Introduction

In this section we return to the origin of tangles: their appearances in graphs.

Tangles in graphs have played a central role in graph minor theory ever since their introduction by Robertson and Seymour in [41]. Formally, a tangle in a graph G is an orientation of all low-order separations of G satisfying certain consistency assumptions. Tangles can be used to locate, and thereby capture the essence of, highly connected substructures in G in that every such substructure defines a tangle in G by orienting each low-order separation of G towards the side containing most or all of that substructure. In view of this, if some tangle in G contains the separation (A, B), we think of A and B as the 'small' and the 'big' side of (A, B) in that tangle, respectively. Our main result will make this intuition concrete.

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

In [17] Diestel raised the question whether all tangles in graphs arise in the above fashion, that is, whether all graph tangles are decided by majority vote by some subset of the vertices:

Problem 4.4.1. 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 [22], where Elbracht showed that such a set X always exists if G is $(k-1)$-connected and has at least $4(k-1)$ vertices. However Elbracht's approach relies heavily on the $(k-1)$-connectedness of the graph and offers no line of attack for the general problem. Finding an answer for arbitrary graphs appears to be hard.

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

In this section, we consider a fractional version of Diestel's question and answer it affirmatively, making precise the notion that B is the 'big' side of a separation $(A, B) \in \tau$: given a k-tangle τ in G, rather than finding a vertex set X which decides τ by majority vote, we find a weight function $w\colon V(G) \to \mathbb{N}$ on the vertices such that for all separations (A, B) of order $< k$ we have $(A, B) \in \tau$ if and only if the vertices in B have higher total weight than those in A.

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

In Section 4.4.2 we will formally define separations and tangles, and 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 4.4.3 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 [28] 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 4.4.3 the fractional version of Problem 4.4.1 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.

4.4.2 Weighted deciders

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

Our main result of this section is the following:

Theorem 30. *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 30 in the remainder of this section. Our general strategy will be as follows: we define a partial order on the separations of G and consider the set of those separations of the k-tangle τ that are maximal in this partial order. For these separations we will be able to show that, on average, their separators divide each other so that they lie more on the 'big' side of each other, where 'big' is the big side according to τ. This will enable us to use a result from linear programming to find a weight function assigning weights to the vertices of these separators such that this weight function decides all these maximal separations of τ correctly. The nature of the partial order will then ensure that this weight function in fact decides all separations in τ correctly.

For a graph G there is a partial order on the separations of G given by letting $(A,B)\leqslant(C,D)$ if and only if $A\subseteq C$ and $B\supseteq D$. One of the main ingredients for the proof of Theorem 30 is the following observation about those separations in a tangle τ that are maximal in τ with respect to this partial order. It says, roughly, that they divide each other's separators so that, on average, those separators lie more on the big side of the separation than on the small side, according to the tangle.

Lemma 4.4.2. *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 4.4.3 (Tucker's Theorem [45])**.** *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 30.

Proof of Theorem 30. Let a finite graph $G = (V, E)$ and a k-tangle τ in G be given. Since G is finite it suffices to find a weight function $w\colon V \to \mathbb{R}_{\geqslant 0}$ such that a separation (A, B) of order $< k$ lies in τ precisely if $w(A) < w(B)$; by the density of the rationals in the reals, this w can then be turned into such a weight function with values in $\mathbb{N}$.

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

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

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

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

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

for $i, j \leqslant n$. Let M be the matrix $\{m_{ij}\}_{i,j \leqslant n}$. Observe that, by Lemma 4.4.2, we have $m_{ij} + m_{ji} > 0$ for all $i \neq j$ and hence the matrix $M + M^T$ has positive entries everywhere but on its diagonal (where it has zeros). We further define

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

Then K is skew-symmetric, that is, $K^T = -K$. Let $x = (x_1, \dots, x_n)^T$ be the vector obtained by applying Lemma 4.4.3 to K. We define a weight function $w\colon V \to \mathbb{R}$ by

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

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

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

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

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

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 such 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 $v \in B_i \smallsetminus A_i$, and increasing the value of $w(v)$ by ε. □

We conclude this section with the remark that Theorem 30 and its proof extend to tangles in hypergraphs without any changes. Even more generally, the following version of Theorem 30 can be established with exactly the same proof as well:

Theorem 31. *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 $w(A) < w(B)$.*

In Theorem 31 a set separation of some ground-set V is a pair (A,B) of subsets of V with $A \cup B = V$. 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 $\vec{U}$. As for graphs, a partial order on the set separations of V is given by letting $(A, B) \leqslant (C, D)$ if $A \cup C$ and $B \supseteq D$.

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 $\vec{U}$. (See [17] for more on the relation between graph tangles and profiles.) Therefore Theorem 31 indeed applies to tangles in graphs and hypergraphs as well.

Theorem 31 holds with the same proof as Theorem 30, since Lemma 4.4.2 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.

4.4.3 Edge-tangles

A related object of study (cf. [20, 36]) 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 32. *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 bipartion (A, B) is defined as $|(A, B)| \coloneqq \sum_{(u,v) \in A \times B} e(u, v)$. 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 32:

Theorem 33. *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 30 and to then transform it into the weighted vertex decider w. So let us first show an analogue of Lemma 4.4.2 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 4.4.2:

Lemma 4.4.4. *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. □

We are now ready to prove Theorem 33:

Proof of Theorem 33. As in the proof of Theorem 30, 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 such 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 $(A_1, B_1), \ldots, (A_n, B_n)$. Just as in Theorem 30 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 4.4.4, $M + M^T$ has positive entries everywhere but on the diagonal, where it is zero. We are now in the same situation as in the proof of Theorem 30 and can find some vector $x \in \mathbb{R}^n_{\geqslant 0}$ such that either $(Mx)_i > 0$ on all i, or x has exactly one non-zero entry, say x_i, and $(Mx)_j > 0$ for all $j \neq i$.

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

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

Note that $\overline{w}$ is symmetric. For the same reason as in Theorem 30, 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

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

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 30, in the latter case we can pick an arbitrary vertex $v \in B_i$ and increase $w(v)$ by some small $\varepsilon > 0$ to achieve $w(B_i) > w(A_i)$ while keeping $w(B_j) > w(A_j)$ for all other maximal elements of τ. $\square$

Remarkably, and in contrast to Theorem 30, Theorem 32 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 k, a *k-edge-tangle* of H is a set τ consisting of exactly one (A, B) or (B, A) for every cut (A, B) of order $< k$, with the additional properties that τ has no subset $\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\}$ such that $B_1 \cap B_2 \cap B_3 = \emptyset$, and that τ contains no cut (A, B) for which B is incident with fewer than k hyperedges of H.

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

Theorem 32 thus asserts that if H is just a (multi-)graph, i.e., if every hyperedge in E has size 2, then every k-edge-tangle of H has such a weighted decider. We are now going to construct an example demonstrating that this may fail for hypergraphs H with hyperedges of size $\geqslant 3$.

Example 4.4.5. For some natural number $k \geqslant 6$ let ℓ be an integer with $3 \leqslant \ell \leqslant \frac{k}{2}$. Let V be the set of all ℓ-element subsets of $[k] = \{1, \ldots, k\}$. Let the set E of hyperedges consist of, for each $i \in [k]$, the set of all $v \in V$ that contain i. Note that each of these k many hyperedges of H has size $\binom{k-1}{\ell-1}$, making H a uniform ℓ-regular hypergraph.

Theorem 34. *Let H be as in Example 4.4.5. Then H has a k-edge-tangle with no weighted decider.*

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

We can therefore define

$$\tau \coloneqq \{(A, B) \in S_k \mid \cup A \subsetneq [k]\} \ .$$

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

To see that τ is a k-edge-tangle we note that by the above observation τ contains exactly one of (A, B) or (B, A) for every cut $(A, B) \in S_k$. Furthermore if $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$ then any element of V containing at least one point each from $[k] \smallsetminus \cup A_1$, from $[k] \smallsetminus \cup A_2$, and from $[k] \smallsetminus \cup A_3$ lies in $B_1 \cap B_2 \cap B_3$, which is hence non-empty since such a $v \in V$ exists by $\ell \geqslant 3$. Finally for each $(A, B) \in \tau$ the set B is incident with each hyperedge of H since $\cup B = [k]$. Thus τ is indeed a k-edge-tangle.

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

$$A_i \coloneqq \{v \in V \mid i \notin v\} \qquad \text{and} \qquad B_i \coloneqq \{v \in V \mid i \in v\} \ ,$$

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

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

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

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

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

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

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

A construction analogous to Example 4.4.5 was found independently by Geelen [28] 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 4.4.5 can also be used to show that allowing weighted deciders to take values in $\mathbb{R}$ rather than $\mathbb{N}$ does not suffice to guarantee their existence for edge-tangles of hypergraphs: for $k = 2\ell$ the tangle described in Theorem 34 has no weighted decider with real-valued and possibly negative weights either, with the same proof.

Chapter 5

Infinite tangle theory

Having completed our tour through finite tangle theory in Chapter 4, let us now extend some of those results to infinite separation systems. Our focus here will be tree-of-tangles theorems as seen in Section 4.1, with ends of graphs as their primary application. Generally speaking the infinite tree-of-tangles theorems in this chapter will be obtained via compactness from finite results.

Our first stop will be a study of ends in graphs and the types of tangles they induce: in Section 5.1 we will see that every end of an infinite graph induces a tangle on the set of all separations of finite order. This set of finite-order separations of a graph can be equipped with a topology in a natural way. The tangles induced by ends, viewed as subsets of this space, might be closed or not. Similarly, the k-tangles induced by the ends of that graph might be closed subsets of that space, or they might not be. We will show that whether or not these tangles are closed depends only on the degree and number of dominating vertices of the end inducing a given tangle. Moreover, such a k-tangle is closed if and only if it is decided by a set of exactly k vertices through majority vote; this is a surprisingly strong property when compared to Section 4.4, where we were only able to show that tangles in finite graphs are decided by *weighted* sets of vertices. Section 5.1 is based on [33].

Following that, in Section 5.2 we will go back to abstract separation systems and study various ways in which one can prove tree-of-tangles theorems in an infinite setting. We will see one straightforward way of extending the splinter theorem 18 to infinite separation systems in Section 5.2.3 using compactness. We will apply this infinite splinter theorem to infinite graphs, obtaining a tree-of-tangles theorem for tangles which are not included in some infinite-order tangle.

Finally we will present another way of extending a finite tree-of-tangles theorem to the infinite setting: this method also uses compactness, but applies this compactness differently than the extension of Theorem 18, resulting in complementary sets of assumptions and a similar but different application in infinite graphs. Our comparison of the two approaches will equip us with the tools necessary to strengthen the application for graphs to include all tangles that are closed in some natural topology. We prove that this is indeed a strengthening by showing that all tangles not included in some infinite-order tangle are indeed closed in that topology. Section 5.2 is joint work with Christian Elbracht and Maximilian Teegen and based on [27], with the exception of Section 5.2.5, which is as-of-yet unpublished work by myself.

5.1 Ends as tangles

5.1.1 Introduction

Our first object of study in infinite tangle theory are tangles in infinite graphs. In [11], it was shown how the set Θ of tangles of infinite order of an arbitrary infinite graph can be used to compactify that graph, much in the same way as the set Ω of ends of a connected locally finite graph can be used to compactify it. Indeed, if a graph G is connected and locally finite, these compactifications $|G|_\Theta$ and $|G|_\Omega$ of G coincide. This is because every end ω of an infinite graph G induces a tangle $\tau = \tau_\omega$ of order $\aleph_0$ in G, and for locally finite connected G the map $\omega \mapsto \tau_\omega$ is a bijection between the set Ω of ends of G and the set Θ of its $\aleph_0$-tangles. (Graphs that are not locally finite have $\aleph_0$-tangles that are not induced by an end.)

In [11], a natural topology on the set $\vec{S} = \vec{S}_{\aleph_0}(G)$ of separations of finite order of G was defined. A tangle τ induced by an end of G is a closed set in this topology if and only if τ is defined by an $\aleph_0$-block in G, that is, if there is an $\aleph_0$-block K in G with $K \subseteq B$ for all separations (A, B) in τ.

Our research expands on this latter result. Every end ω of a graph induces not only a tangle of infinite order in G, but for each $k \in \mathbb{N}$ the end ω induces a k-tangle in G. The set $\vec{S}_k$ of all separations (A, B) of G with $|A \cap B| < k$ is a closed set in $\vec{S}$, and thus if the tangle τ induced by ω in G is a closed set in $\vec{S}$, the k-tangle $\tau \cap \vec{S}_k$ induced by ω will be closed as well. However, it is possible that a tangle τ in G of infinite order fails to be closed in $\vec{S}$, while its restrictions $\tau \cap \vec{S}_k$ to $\vec{S}_k$ are closed for some, or even all, $k \in \mathbb{N}$. In this section we characterize the ends of G by the behaviour of their tangles, as follows: We show that, for an end ω and its induced tangle τ, the restriction $\tau \cap \vec{S}_k$ to $\vec{S}_k$ is a closed set in $\vec{S}$ if and only if

$$\deg(\omega) + \mathrm{dom}(\omega) \geqslant k,$$

where $\deg(\omega)$ and $\mathrm{dom}(\omega)$ denote the vertex degree and number of vertices dominating ω, respectively.

We further show that τ is closed in $\vec{S}$ if and only if ω is dominated by infinitely many vertices.

A question raised in [17] asks whether for a k-tangle τ in a finite graph G one can always find a set X of vertices which decides τ by majority vote, in the sense that $(A, B) \in \tau$ if and only if $|A \cap X| < |B \cap X|$, for all $(A, B) \in \vec{S}_k$. This problem is still open in general, although some process has been made recently (see Section 4.4 and [22, 25]). We establish an analogue in the infinite setting: we show that for an end ω of G and its induced k-tangle $\tau \cap \vec{S}_k$ in G, the existence of a finite set X which decides $\tau \cap \vec{S}_k$ in the above sense is equivalent to $\tau \cap \vec{S}_k$ being a closed set in $\vec{S}$.

This section is organized as follows: Section 5.1.2 contains the basic definitions and some notation. Following that, in Section 5.1.3, we recall the core concepts and results from [11] that are relevant to our studies, including the topology defined on $\vec{S}$. Finally, in Section 5.1.4, we prove our main results Theorem 36 and Theorem 37. The first of these characterises the ends of a graph by the behaviour of their tangles, and the second shows that $\tau \cap \vec{S}_k$ being a closed set in $\vec{S}$ for a k-tangle τ induced by some end ω of G is equivalent to both $\deg(\omega) + \mathrm{dom}(\omega) \geqslant k$ and to $\tau \cap \vec{S}_k$ being decided by some finite set of vertices.

5.1.2 Separations, tangles, and their topology

Throughout this section $G = (V, E)$ will be a fixed infinite graph. Let us recall the relevant definitions for tangles in graphs, and their extensions to infinite graphs.

The *order* of a separation (A, B) or $\{A, B\}$ of G is the cardinality $|A \cap B|$ of its separator. For a cardinal κ we write $S_\kappa = S_\kappa(G)$ for the set of all unoriented separations of G of order $< \kappa$. If S is a set of unoriented separations we write $\vec{S}$ for the corresponding set of oriented separations, that is, the set of all separations (A, B) with $\{A, B\} \in S$. Consequently we write $\vec{S}_\kappa$ for the set of all separations (A, B) of G with $|A \cap B| < \kappa$.

If S is a set of unoriented separations of G, an *orientation* of S is a set $O \subseteq \vec{S}$ such that O contains precisely one of (A, B) or (B, A) for every $\{A, B\} \in S$. A *tangle of S in G* is an orientation τ of S such that there are no $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$ for which $G[A_1] \cup G[A_2] \cup G[A_3] = G$.

Properties of sets of separations of finite graphs, including their tangles, often generalize to sets of separations of infinite graphs but not always. Those sets of separations to which these properties tend to generalize can be identified, however: they are the sets of separations that are closed in a certain natural topology [18]. Let us define this topology next. It is analogous to the topology of a profinite abstract separation system defined in [18].[1]

From here on we denote by $S = S_{\aleph_0}(G)$ the set of all (unoriented) separations of G of finite order. Thus $\vec{S} = \vec{S}_{\aleph_0}(G)$ is the set of all separations (A, B) of G with $A \cap B$ finite.

We define our topology on $\vec{S}$ by giving it the following basic open sets. Pick a finite set $Z \subseteq V$ and an oriented separation (A_Z, B_Z) of $G[Z]$. Then declare as open the set $O(A_Z, B_Z)$ of all $(A, B) \in \vec{S}$ such that $A \cap Z = A_Z$ and $B \cap Z = B_Z$. We shall say that these (A, B) *induce* (A_Z, B_Z) on Z, writing $(A_Z, B_Z) =: (A, B)\restriction Z$, and that (A, B) and (A', B') *agree on* Z if $(A, B)\restriction Z = (A', B')\restriction Z$.

It is easy to see that the sets $O(A_Z, B_Z)$ do indeed form the basis of a topology on $\vec{S}$. Indeed, $(A, B) \in \vec{S}$ induces (A_1, B_1) on Z_1 and (A_2, B_2) on Z_2 if and only if it induces on $Z = Z_1 \cup Z_2$ some separation (A_Z, B_Z) which in turn induces (A_i, B_i) on Z_i for both i. Hence $O(A_1, B_1) \cap O(A_2, B_2)$ is the union of all these $O(A_Z, B_Z)$.

As we shall see, the intuitive property of tangles in finite graphs that they describe, if indirectly, some highly cohesive region of that graph – however 'fuzzy' this may be in terms of concrete vertices and edges – will extend precisely to those tangles of S that are closed in $\vec{S}$.

5.1.3 End tangles of S

We think of an oriented separation $(A, B) \in \vec{S}$ as pointing towards B, or being oriented towards B. In the same spirit, given an end ω of G, we say that (A, B) *points towards* ω, and that ω *lives in* B, if some (equivalently: every) ray of ω has a tail in B. Furthermore, if (A, B) points to an end ω, then (B, A) *points away from* ω.

[1] Even though $\vec{S}$ itself is not usually profinite in the sense of [18], the topology we define on $\vec{S}$ is the subspace topology of $\vec{S}$ as a subspace of the (profinite) system of all oriented separations of G, equipped with the inverse limit topology from [18].

Clearly, for every end of G and every $\{A, B\} \in S$, precisely one orientation of $\{A, B\}$ points towards that end. In this way, every end ω of G defines an orientation of S by orienting each separation in S towards ω:

$$\tau = \tau_\omega := \{(A, B) \in \vec{S} \mid \text{every ray of } \omega \text{ has a tail in } B\}$$

It is easy to see ([11]) that this is a tangle in G. We call it the *end tangle induced* on S by the end ω.

Note that every end tangle contains all separations of the form (A, V) for finite $A \subseteq V$, and thus no separation of the form (V, B). Furthermore, any two ends induce different end tangles. Our aim in this section is to recall from [11] some properties of the end-tangles of S that we shall later extend to its subsets S_k. For the convenience of the reader, and also in order to correct an inessential but confusing error in [11], we repeat some of the material from [11] here to make our presentation self-contained.

Let us first see an example of an end tangle that is not closed in $\vec{S}$.

Example 5.1.1. *If G is a single ray $v_0v_1\dots$ with end ω, say, then $\tau = \tau_\omega$ is not closed in $\vec{S}$.*

Indeed, τ contains $(\emptyset, V)$, and hence does not contain $(V, \emptyset)$. But for every finite $Z \subseteq V$ the restriction $(Z, \emptyset)$ of $(V, \emptyset)$ to Z is also induced by the separation $(\{v_0, \dots, v_n\}, \{v_n, v_{n+1}, \dots\}) \in \tau$ for every n large enough that $Z \subseteq \{v_0, \dots, v_{n-1}\}$. So $(V, \emptyset) \in \vec{S} \smallsetminus \tau$ has no open neighbourhood in $\vec{S} \smallsetminus \tau$.□

Here is an example of an end tangle that is closed in $\vec{S}$. Unlike our previous example, it describes a highly cohesive part of G.

Example 5.1.2. *If $K \subseteq V$ spans an infinite complete graph in G, then*

$$\tau = \{\, (A, B) \in \vec{S} \mid K \subseteq B \,\} \tag{1}$$

is a closed set in $\vec{S}$.

We omit the easy proof. But note that τ is indeed an end tangle: it is induced by the unique end of G which contains all the rays in K.

Perhaps surprisingly, it is not hard to characterize the end tangles that are closed. They are all essentially like Example 5.1.2: we just have to generalize the infinite complete subgraph used appropriately. Of the two obvious generalizations, infinite complete minors [43] or subdivisions of infinite complete graphs [42], the latter turns out to be the right one.

Let κ be any cardinal. A set of at least κ vertices of G is $(< \kappa)$-*inseparable* if no twoof them can be separated in G by fewer than κ vertices. A maximal $(< \kappa)$-inseparable set of vertices is a κ-*block*. For example, the branch vertices of a TK_κ are $(< \kappa)$-inseparable. Conversely:

Lemma 5.1.3. *When κ is infinite, every $(< \kappa)$-inseparable set of vertices in G contains the branch vertices of some $TK_\kappa \subseteq G$.*

Proof. Let $K \subseteq V$ be $(< \kappa)$-inseparable. Viewing κ as an ordinal we can find, inductively for all $\alpha < \kappa$, distinct vertices $v_\alpha \in K$ and internally disjoint v_α-v_β paths in G for all $\beta < \alpha$ that also have no inner vertices among those v_β or on any of the paths chosen earlier; this is because $|K| \geqslant \kappa$, and no two vertices of K can be separated in G by the $< \kappa$ vertices used up to that time. □

The original statement of Lemma 5.1.3 in [11, Lemma 5.4] asserted that for infinite κ a set $K \subseteq V$ is a κ-block in G if and only if it is the set of branch vertices of some $TK_\kappa \subseteq G$. It turns out that both directions of that assertion were false: the set of branch vertices of a $TK_\kappa \subseteq G$ is certainly $(< \kappa)$-inseparable, but might not be maximal with this property and hence not a κ-block. Conversely, if K is a κ-block, there might not be a $TK_\kappa \subseteq G$ whose set of branch vertices is precisely K: if $|K| > \kappa$ this is certainly not possible, but even if $|K| = \kappa$ one might not be able to find a TK_κ in G whose branch vertices are all of K. If for instance the graph G is a clique on κ vertices that is missing exactly one edge, then $K = V(G)$ is a κ-block in G but not the set of branch vertices of a $TK_\kappa \subseteq G$.

For the main result of this section we need one more observation:

Lemma 5.1.4. *If τ is an end tangle of G and $(A,B),(C,D) \in \tau$ then $(A \cup C\,,\, B \cap D) \in \tau$.*

Proof. Observe first that $(A \cup C\,,\, B \cap D)$ is a separation of G with finite order and thus lies in $\vec{S}$. Moreover, if a ray of G has a tail in B and a tail in D, then that ray also has a tail in $B \cap D$, from which it follows that $(A \cup C\,,\, B \cap D) \in \tau$, as claimed. $\square$

We can now prove our main result of this section. Let us say that a set $K \subseteq V$ *defines* an end tangle τ if τ satisfies (1).

Theorem 35 ([11])**.** *Let G be any graph. An end tangle of G is closed in $\vec{S}$ if and only if it is defined by an $\aleph_0$-block.*

Proof. Suppose first that τ is an end tangle that is defined by an $\aleph_0$-block K. To show that τ is closed, we have to find for every $(A,B) \in \vec{S} \smallsetminus \tau$ a finite set $Z \subseteq V$ such that no $(A',B') \in \vec{S}$ that agrees with (A,B) on Z lies in τ. Since $(A,B) \notin \tau$, we have $K \subseteq A$; pick $z \in K \smallsetminus B$. Then every $(A',B') \in \vec{S}$ that agrees with (A,B) on $Z := \{z\}$ also also lies in $\vec{S} \smallsetminus \tau$, since $z \in A' \smallsetminus B'$ and this implies $K \nsubseteq B'$.

Conversely, consider any end tangle τ and let

$$K := \bigcap \{\, B \mid (A,B) \in \tau \,\}.$$

No two vertices in K can be separated by in G by a finite-order separation: one orientation (A,B) of this separation would be in τ, which would contradict the definition of K since $A \smallsetminus B$ also meets K. If K is infinite, it will clearly be maximal with this property, and hence be an $\aleph_0$-block. This $\aleph_0$-block K will define τ: by definition of K we have $K \subseteq B$ for ever $(A,B) \in \tau$, while also every $(A,B) \in \vec{S}$ with $K \subseteq B$ must be in τ: otherwise $(B,A) \in \tau$ and hence $K \subseteq A$ by definition of K, but $K \nsubseteq A \cap B$ because this is finite. Hence τ will be defined by an $\aleph_0$-block, as desired for the forward implication.[2]

It thus suffices to show that if K is finite then τ is not closed in $\vec{S}$, which we shall do next.

Assume that K is finite. We have to find some $(A,B) \in \vec{S} \smallsetminus \tau$ that is a limit point of τ, i.e., which agrees on every finite $Z \subseteq V$ with some $(A',B') \in \tau$. We choose $(A,B) := (V,K)$, which lies in $\vec{S} \smallsetminus \tau$ since $(K,V) \in \tau$.

[2] Whether or not τ is closed in $\vec{S}$ is immaterial; we just did not use this assumption.

To complete our proof as outlined, let any finite set $Z \subseteq V$ be given. For every $z \in Z \smallsetminus K$ choose $(A_z, B_z) \in \tau$ with $z \in A_z \smallsetminus B_z$: this exists, because $z \notin K$. Since $(K, V) \in \tau$, by Lemma 5.1.4 we have $(A', B') \in \tau$ for

$$A' := K \cup \bigcup_{z \in Z \smallsetminus K} A_z \quad \text{and} \quad B' := V \cap \bigcap_{z \in Z \smallsetminus K} B_z \,.$$

As desired, $(A', B') \restriction Z = (A, B) \restriction Z$ (which is $(Z, Z \cap K)$, since $(A, B) = (V, K)$): every $z \in Z \smallsetminus K$ lies in some A_z and outside that B_z, so $z \in A' \smallsetminus B'$, while every $z \in Z \cap K$ lies in $K \subseteq A'$ and also, by definition of K, in every B_z (and hence in B'), since $(A_z, B_z) \in \tau$. □

This proof of Theorem 35 concludes our exposition of material from [11].

We conclude the section with the remark that an end tangle τ of an infinite graph G is defined by an $\aleph_0$-block if and only if it is defined by some set $X \subseteq V$: for if τ is defined by some $X \subseteq V$, then X must be infinite, as otherwise $\{X, V\} \in S$ would witness that X does not define τ. But if X is infinite then so is $K := \bigcap \{B \mid (A, B) \in \tau\} \supseteq X$, and we can follow the proof of Theorem 35 to show that K is an $\aleph_0$-block defining τ.

5.1.4 End tangles in S_k

For $k \in \mathbb{N}$ we write S_k for the subset of S containing all separations of G of order $< k$. In this section we shall extend Theorem 35 from the previous section to these S_k, where we will also be able to prove a stronger and wider statement which will allow us to classify ends by the behaviour of their end tangles – specifically, by the set of those $k \in \mathbb{N}$ for which their end tangle is closed in $\vec{S}_k$.

For every $k \in \mathbb{N}$ the set $\vec{S}_k$ is a closed set in $\vec{S}$. Given an end ω of G with end tangle $\tau = \tau_\omega$, the *tangle induced by ω in S_k* is $\tau \cap \vec{S}_k$. We say that ω *induces a closed tangle in S_k* if $\tau \cap \vec{S}_k$ is a closed set in $\vec{S}_k$, where the latter is equipped with the subspace topology of $\vec{S}$.

We seek to classify and characterize ends by the set of k for which their induced tangles are closed in S_k, similarly to Theorem 35. For $\ell < k$ the set $\vec{S}_\ell$ is closed in $\vec{S}_k$, and hence any end inducing a closed tangle in S_k also induces a closed tangle in S_ℓ. Therefore, the ends of G all fall into one of the following three categories:

(i) Ends inducing a closed tangle in S;

(ii) ends whose induced tangle in S is not closed in $\vec{S}$, but whose induced tangle in S_k is closed in $\vec{S}_k$ for every $k \in \mathbb{N}$;

(iii) ends whose induced tangle in S_k is not closed in $\vec{S}_k$ for some $k \in \mathbb{N}$.

The ends belonging to the first category are characterized in Theorem 35: they are those ends whose tangle in $\vec{S}_{\aleph_0}$ is defined by an $\aleph_0$-block. In the remainder of this section we shall characterize the ends which fall into the latter two categories. Furthermore, for an end ω belonging to the third category, there exists a least $k \in \mathbb{N}$ for which the tangle induced by ω in S_k is not closed in $\vec{S}_k$. We shall determine this k and show that it depends just on the vertex degree and number of dominating vertices of ω.

Let us see examples of ends belonging to the third and second category, respectively:

Example 5.1.5. Let G be as in Example 5.1.1, that is, a single ray $v_0v_1\dots$ with end ω. The same argument as in Example 5.1.1 shows that $\tau \cap \vec{S}_k$ is not closed in $\vec{S}_k$ for $k \geqslant 2$. However, ω does induce a closed tangle in $\vec{S}_1$: the set $\tau \cap \vec{S}_1 = \{(\emptyset, V)\}$ is closed in $\vec{S}_1$.

Example 5.1.6. Let G be the infinite grid, ω the unique end of G and τ the tangle induced by ω in $S = S(G)$. Since G is locally finite it does not contain an $\aleph_0$-block. Therefore τ cannot be defined by an $\aleph_0$-block and is thus not closed in $\vec{S}$ by Theorem 35. However, for every $k \in \mathbb{N}$, it is easy to see that $\tau \cap \vec{S}_k$ is closed in $\vec{S}_k$: indeed, for fixed $k \in \mathbb{N}$, the size of A is bounded in terms of k for every $(A, B) \in \tau \cap \vec{S}_k$. Thus, for every $(A', B') \in \vec{S}_k \setminus \tau$, any $Z \subseteq V(G)$ with $|Z \cap A'|$ sufficiently large witnesses that (A', B') does not lie in the closure of τ: if Z is large enough that $|Z \cap A'| > |A|$ for every $(A, B) \in \tau \cap \vec{S}_k$, then no $(A, B) \in \tau \cap \vec{S}_k$ will agree with (A', B') on Z.

Example 5.1.6 shows that the end tangle τ of the infinite grid is not defined by a set $X \subseteq V$ in the sense of (1). Moreover, even for any fixed $k \in \mathbb{N}_{\geqslant 5}$, we cannot find a set $X \subseteq V$ with $X \subseteq B$ for all $(A, B) \in \tau \cap \vec{S}_k$: for every $x \in X$ the separation $(\{x\} \cup N(x)\,, V \setminus \{x\})$ lies in $\tau \cap \vec{S}_k$. Thus, even the tangles that τ induces in $\vec{S}_k$ are not defined by any subset of V. However, they come reasonably close to it: for $X = V$ and any $(A, B) \in \tau$, the majority of X lies in B, that is, we have $|A \cap X| < |B \cap X|$. In fact, for any fixed $k \in \mathbb{N}$, any finite set $X \subseteq V$ that is at least twice as large as $\max\{|A| \mid (A, B) \in \tau \cap \vec{S}_k\}$ has the property that $|A \cap X| < |B \cap X|$ for every $(A, B) \in \tau \cap \vec{S}_k$. Therefore, even though no $\tau \cap \vec{S}_k$ is defined by a $X \subseteq V$, we can for each $k \in \mathbb{N}$ find a (finite) $X \subseteq V$ which 'defines' $\tau \cap \vec{S}_k$ by simple majority.

To make the above observation formal, for an end tangle τ of G, let us call a set $X \subseteq V$ a *decider set* for τ (resp., for $\tau \cap \vec{S}_k$) if we have $|A \cap X| < |B \cap X|$ for every $(A, B) \in \tau$ (resp., $(A, B) \in \tau \cap \vec{S}_k$). Thus, if we have a decider set for an end tangle τ, given a separation $(A, B) \in \vec{S}$, the decider set tells us which of (A, B) and (B, A) lies in τ by a simple majority vote. By Theorem 35, every end tangle that is closed in $\vec{S}$ has a decider set, since the $\aleph_0$-block defining it is such a decider set. In analogy with this, we shall show in the remainder of this section that for an end tangle τ its induced tangle $\tau \cap \vec{S}_k$ is closed in $\vec{S}_k$ if and only if $\tau \cap \vec{S}_k$ has a finite decider set. Such a finite decider set can be thought of as a local encoding of the tangle, or a local witness to the tangle being closed in $\vec{S}_k$.

In contrast, it is easy to see that every end tangle has an infinite decider set:

Proposition 5.1.7. *For any end ω of G the end tangle τ induced by ω has an infinite decider set.*

Proof. For any ray $R \in \omega$ its vertex set $V(R)$ is a decider set for τ. □

Since no end tangle of S can have a finite decider set the existence of decider sets for end tangles is thus only interesting for finite decider sets of the tangle's restrictions to some $\vec{S}_k$.

We shall complement this local witness of a given end tangle being closed with a more global type of witness. For this we need to introduce some notation.

The *(vertex) degree* $\deg(\omega)$ of an end ω of G is the largest size of a family of pairwise disjoint ω-rays[3]. A vertex $v \in V$ *dominates* an end ω if it sends infinitely

[3]Here our notation deviates from that in [10], where $d(\omega)$ is used for the degree of ω.

many disjoint paths to some (equivalently: to each) ray in ω. We write $\mathrm{dom}(\omega)$ for the number of vertices of G which dominate ω. An end ω is *undominated* if $\mathrm{dom}(\omega) = 0$; it is *finitely dominated* if finitely many (including zero) vertices of G dominate ω; and, finally, ω is *infinitely dominated* if $\mathrm{dom}(\omega) = \infty$. We will show that the category that an end ω belongs to depends just on these parameters $\deg(\omega)$ and $\mathrm{dom}(\omega)$. Concretely, we will show the following:

Theorem 36. *Let τ the end tangle induced by an end ω of G. Then the following statements hold:*

(i) *τ is closed in $\vec{S}$ if and only if $\mathrm{dom}(\omega) = \infty$.*

(ii) *τ is not closed in $\vec{S}$ but $\tau \cap \vec{S}_k$ is closed in $\vec{S}_k$ for every $k \in \mathbb{N}$ if and only if $\deg(\omega) = \infty$ and $\mathrm{dom}(\omega) < \infty$.*

(iii) *$\tau \cap \vec{S}_k$ is not closed in $\vec{S}_k$ for some $k \in \mathbb{N}$ if and only if $\deg(\omega) + \mathrm{dom}(\omega) < \infty$.*

Theorem 36 will be a consequence of Theorem 35 and the following theorem, which characterizes for which $k \in \mathbb{N}$ a given end tangle is closed in $\vec{S}_k$ and makes the connection to finite decider sets:

Theorem 37. *Let τ be the end tangle induced by an end ω of G and let $k \in \mathbb{N}$. Then the following are equivalent:*

(i) *$\tau \cap \vec{S}_k$ is closed in $\vec{S}_k$;*

(ii) *$\deg(\omega) + \mathrm{dom}(\omega) \geqslant k$;*

(iii) *$\tau \cap \vec{S}_k$ has a finite decider set;*

(iv) *$\tau \cap \vec{S}_k$ has a decider set of size exactly k.*

Let us first prove Theorem 36 from Theorem 35 and Theorem 37:

Proof of Theorem 36. We will show (i) using Theorem 35 and (ii) with Theorem 37. The third statement is then an immediate consequence of the first two and the fact that an end tangle which is closed in $\vec{S}$ is also closed in $\vec{S}_k$.

To see that (i) holds, let us first suppose that τ is closed in $\vec{S}$. Then, by Theorem 35, τ is defined by an $\aleph_0$-block $K \subseteq V$. It is easy to see that every vertex in K dominates ω, hence $\mathrm{dom}(\omega) = \infty$ as K is infinite.

For the converse, suppose that ω is infinitely dominated, and let us show that τ is closed in $\vec{S}$. For every separation $(A, B) \in \tau$ every vertex dominating ω must lie in B. Therefore $K := \bigcap\{B \mid (A, B) \in \tau\}$ is infinite and thus, as seen in the proof of Theorem 35, an $\aleph_0$-block defining τ, which is hence closed in $\vec{S}$.

Let us now show (ii). By (i), τ is not closed in $\vec{S}$ if and only if $\mathrm{dom}(\omega) < \infty$. On the other hand, by Theorem 37, $\tau \cap \vec{S}_k$ is closed in $\vec{S}_k$ for all k if and only if $\deg(\omega) + \mathrm{dom}(\omega) \geqslant k$ for all $k \in \mathbb{N}$, that is, if at least one summand is infinite. Thus (ii) holds.

Claim (iii) now follows either from claim (i) and (ii), or directly from Theorem 37. □

We will conclude this section by proving Theorem 37. For this we shall need two lemmas. The first lemma can be seen as an analogue of Menger's Theorem between a vertex set and an end. Given a set $X \subseteq V$ and an end ω, we say that $F \subseteq V$ *separates X from ω* if every ω-ray which meets X also meets F.

Lemma 5.1.8. *Let ω be an undominated end of G and $X \subseteq V$ a finite set. The largest size of a family of disjoint ω-rays which start in X is equal to the smallest size of a set $T \subseteq V$ separating X from ω.*

Proof. Let T be a set separating X from ω of minimal size. Clearly, a family of disjoint ω-rays which all start in X cannot be larger than T, as every ray in that family must meet T. So let us show that we can find a family of $|T|$ disjoint ω-rays starting in X.

Observe that, as ω is undominated, for every vertex $v \in V$ we can find a finite set $T_v \subseteq V \smallsetminus \{v\}$ which separates v from ω. Thus, for every finite set $Y \subseteq V$ we can find a finite set in $V \smallsetminus Y$ separating Y from ω: for instance, the set $\bigcup\{T_v \smallsetminus Y \mid v \in Y\}$.

Pick a sequence of finite sets $T_n \subseteq V$ inductively by setting $T_0 := T$ and picking as T_n a set of minimal size with the property that $T_{n-1} \cap T_n = \emptyset$ and that T_n separates T_{n-1} from ω; these sets exist by the above observation. Let C_n be the component of $G - T_n$ that contains ω. Clearly $T_n \subseteq C_{n-1}$.

We claim that $C := \bigcap_{n \in \mathbb{N}} C_n = \emptyset$. To see this, consider any $v \in C$ and a shortest v–X path in G. This path must pass through T_n for every $n \in \mathbb{N}$, which is impossible since the separators T_n are pairwise disjoint. Therefore C must be empty.

By the minimality of each T_n, the sets T_n are of non-decreasing size, and furthermore Menger's Theorem yields a family of $|T_n|$ many disjoint paths between T_n and T_{n+1} for each $n \in \mathbb{N}$, as well as $|T_0|$ many disjoint paths between X and T_0. By concatenating these paths we obtain a family of $|T_0| = |T|$ many rays starting in X. To finish the proof we just need to show that these rays belong to ω. To see this, let ω' be another end of G, and T' a finite set separating ω and ω'. Since $C = \emptyset$ we have $C_n \cap T' = \emptyset$ for sufficiently large n, which shows that the rays constructed do not belong to ω' and hence concludes the proof. $\square$

An immediate consequence of Lemma 5.1.8 is that, for an undominated end ω, every finite set $X \subseteq V$ can be separated from ω by at most $\deg(\omega)$ many vertices. In fact, we can state a slightly more general corollary:

Corollary 5.1.9. *Let ω be a finitely dominated end of G and $X \subseteq V$ a finite set. Then X can be separated from ω by some $T \subseteq V$ with $|T| \leqslant \deg(\omega) + \operatorname{dom}(\omega)$.*

Proof. Let D be the set of vertices dominating ω and consider the graph $G' := G - D$ and the set $X' := X \smallsetminus D$. By Lemma 5.1.8 there is a set $T' \subseteq V(G')$ of size at most $\deg(\omega)$ separating X' from ω in G'. Set $T := T' \cup D$. Then T separates X from ω in G and has size $|T| = |T'| + |D| \leqslant \deg(\omega) + \operatorname{dom}(\omega)$. $\square$

The second lemma we shall need for our proof of Theorem 37 roughly states that, for an end of high degree, we can find a large family of disjoint ω-rays whose set of starting vertices is highly connected in G, even after removing the tails of these rays:

Lemma 5.1.10. *Let ω be an end of G and $k \leqslant \deg(\omega) + \operatorname{dom}(\omega)$. Then there are a set $X \subseteq V$ of k vertices and a set $\mathcal{R}$ of disjoint ω-rays with the following properties: every vertex in X is either the start-vertex of a ray in $\mathcal{R}$, or dominates ω and does not lie on any $R \in \mathcal{R}$, and furthermore for any two sets $A, B \subseteq X$ there are $\min(|A|, |B|)$ many disjoint A–B-paths in G whose internal vertices meet no ray in $\mathcal{R}$ and no vertex of X.*

Proof. Pick a set D of vertices dominating ω and a set $\mathcal{R}$ of disjoint ω-rays not meeting D such that $|D| + |\mathcal{R}| = k$; we shall find suitable tails of the rays in $\mathcal{R}$ such that their starting vertices together with D are the desired set X.

Using the fact that the vertices in D dominate ω and that the rays in $\mathcal{R}$ belong to ω, we can pick for each pair x_1, x_2 of elements of $D \cup \mathcal{R}$ an x_1–x_2-path in G in such a way that these paths are pairwise disjoint with the exception of possibly having a common end-vertex in D. Let $\mathcal{P}$ be the set of these paths. Now for each ray in $\mathcal{R}$ pick a tail of that ray which avoids all the paths in $\mathcal{P}$. Let $\mathcal{R}'$ be the set of these tails and X the union of their starting vertices and D. We claim that X and $\mathcal{R}'$ are as desired.

To see this, let us show that for any sets $A, B \subseteq X$ we can find $\min(|A|, |B|)$ many disjoint A–B-paths in G whose internal vertices avoid D as well as $V(R')$ for every $R' \in \mathcal{R}'$. Clearly it suffices to show this for disjoint sets A, B of equal size. So let $A, B \subseteq X$ be two disjoint sets with $n := |A| = |B|$ and let $\mathcal{R}_{A,B}$ be the set of all rays in $\mathcal{R}$ that contain a vertex from A or B. For each pair $(a, b) \in A \times B$ there is a unique path $P \in \mathcal{P}$ such that each of its end-vertices either is a or b (if $a \in D$ or $b \in D$) or lies on a ray in $\mathcal{R}_{A,B}$ which contains a or b; let $P_{a,b}$ be the a–b-path obtained from P by extending it, for each of its end-vertices that is not either a or b, along the corresponding ray in $\mathcal{R}$ up to a or b. Let $\mathcal{P}_{A,B}$ be the set of all these paths $P_{a,b}$. Note that the internal vertices of each path $P_{a,b} \in \mathcal{P}_{A,B}$ meet none of the rays in $\mathcal{R}'$ or vertices in X.

We claim that A and B cannot be separated by fewer than $n = |A|$ vertices in

$$G' := \bigcup_{P_{a,b} \in \mathcal{P}_{A,B}} P_{a,b}\,;$$

the claim will then follow from Menger's Theorem. So suppose that some set $T \subseteq V(G')$ of size less than n is given. Let x and y be the number of vertices in A and B, respectively, whose ray in $\mathcal{R}_{A,B}$ does not meet T. There are xy many paths in $\mathcal{P}_{A,B}$ between these vertices in A and B. Since these paths are disjoint outside their corresponding ray segments, each vertex of T can lie on at most one of them. Thus if $xy \geqslant n$ there must be a T-avoiding path in $\mathcal{P}_{A,B}$ whose end-vertices' rays in $\mathcal{R}_{A,B}$ also do not meet T.

Since $x + y \geqslant n + 1$ we have $xy \geqslant x(n + 1 - x)$. The right-hand side of this inequality, as a function of x with domain $[n - 1]$, is minimized by taking $x = 1$, wherefore it evaluates to n. Thus $xy \geqslant n$, which shows that T does not separate A and B in G'.

We can thus apply Menger's Theorem to obtain n disjoint A–B-paths in G', which are the desired disjoint paths in G whose internal vertices avoid the rays in $\mathcal{R}'$ and vertices in X: the only vertices that are contained both in $V(G')$ as well as in either X or a ray from $\mathcal{R}'$ are vertices from A or B, which cannot be internal vertices of the $n = |A| = |B|$ disjoint A–B-paths. □

We are now ready to prove Theorem 37:

Theorem 37. *Let τ be the end tangle induced by an end ω of G and let $k \in \mathbb{N}$. Then the following are equivalent:*

(i) $\tau \cap \vec{S}_k$ *is closed in* $\vec{S}_k$*;*

(ii) $\deg(\omega) + \operatorname{dom}(\omega) \geqslant k$*;*

(iii) $\tau \cap \vec{S}_k$ *has a finite decider set;*

(iv) $\tau \cap \vec{S}_k$ *has a decider set of size exactly k.*

Proof. We will show (i) $\Rightarrow$ (ii) $\Rightarrow$ (iv) $\Rightarrow$ (iii) $\Rightarrow$ (i).

To see that (i) $\Rightarrow$ (ii), let us first suppose that $\deg(\omega) + \mathrm{dom}(\omega) < k$ and show that $\tau \cap \vec{S}_k$ is not closed in $\vec{S}_k$. Let D be the set of dominating vertices of ω. Since $|D| = \mathrm{dom}(\omega) < k$ the separation (V, D) lies in $\vec{S}_k$. By definition of τ we have $(D, V) \in \tau$ and $(V, D) \notin \tau$. Thus it suffices to show that (V, D) lies in the closure of $\tau \cap \vec{S}_k$ in $\vec{S}_k$. This will be the case if for every finite set $X \subseteq V$ there is a separation $(A, B) \in \tau \cap \vec{S}_k$ with $(A, B) \restriction X = (V, D) \restriction X$.

So let X be a finite subset of V. By Corollary 5.1.9 some set T of size at most $\deg(\omega) + \mathrm{dom}(\omega)$ separates X from ω. Let C be the component of G containing ω. We define the separation (A, B) by setting $A := V \setminus C$ and $B := T \cup C$. Then (A, B) is a separation of G with

$$|(A, B)| = |T| \leqslant \deg(\omega) + \mathrm{dom}(\omega) < k,$$

so $(A, B) \in \vec{S}_k$. In fact (A, B) lies in $\tau \cap \vec{S}_k$ as ω lives in B. Furthermore we have $X \subseteq A$ and $D \subseteq B$ since no vertex dominating ω can be separated from ω by T. Therefore $(A, B) \restriction X = (X, X \cap D) = (V, D) \restriction X$, showing that (V, D) lies in the closure of $\tau \cap \vec{S}_k$ in $\vec{S}_k$.

Let us now show that (ii) $\Rightarrow$ (iv). So let us assume that $\deg(\omega) + \mathrm{dom}(\omega) \geqslant k$. Then by Lemma 5.1.10 we find a set $X \subseteq V$ of size k and a family $\mathcal{R}$ of ω-rays such that every vertex of X either dominates ω or is the start-vertex of a ray in $\mathcal{R}$, and such that for any $A, B \subseteq X$ we can find $\min(|A|, |B|)$ many disjoint A–B-paths in G whose internal vertices meet neither X nor any ray in $\mathcal{R}$.

We claim that X is the desired decider set for $\tau \cap \vec{S}_k$. To see this, let (A, B) be any separation in $\tau \cap \vec{S}_k$; we need to show that $|A \cap X| < |B \cap X|$. Let us write $X_{A \setminus B} := (A \setminus B) \cap X$ and $X_{B \setminus A} := (B \setminus A) \cap X$ as well as $X_{A \cap B} := (A \cap B) \cap X$. It then suffices to prove $|X_{A \setminus B}| < |X_{B \setminus A}|$.

So suppose to the contrary that $|X_{A \setminus B}| \geqslant |X_{B \setminus A}|$. Note first that no vertex in $X_{A \setminus B}$ dominates ω as witnessed by the finite-order separation $(A, B) \in \tau \subseteq S_k$. Therefore, for every vertex in $X_{A \setminus B}$, we have a ray in $\mathcal{R}$ starting at that vertex. Each of those disjoint rays must pass through the separator $A \cap B$, and none of them hits $X_{A \cap B}$. Furthermore by Lemma 5.1.10 there are $|X_{B \setminus A}|$ many disjoint $X_{A \setminus B}$–$X_{B \setminus A}$-paths whose internal vertices avoid $\mathcal{R}$ and X. These paths, too, must pass the separator $A \cap B$ without meeting $X_{A \cap B}$ or any of the rays above. Thus we have

$$|A \cap B| \geqslant |X_{A \cap B}| + |X_{A \setminus B}| + |X_{B \setminus A}| = |X| = k,$$

a contradiction since $(A, B) \in S_k$ and hence $|A \cap B| < k$. Therefore we must have $|X_{A \setminus B}| < |X_{B \setminus A}|$, which immediately implies $|A \cap X| < |B \cap X|$.

Finally, let us show that (iii) $\Rightarrow$ (i). So let $X \subseteq V(G)$ be a finite decider set for $\tau \cap \vec{S}_k$. We need to show that no $(A, B) \in S_k \setminus \tau$ lies in the closure of $\tau \cap \vec{S}_k$. For this let $(A, B) \in S_k \setminus \tau$ be given; then X witnesses that (A, B) does not lie in the closure of τ. To see this, let any $(C, D) \in \tau$ be given. Since X is a decider set for τ we have $|C \cap X| < |D \cap X|$, and since $(A, B) \notin \tau$ we have $|A \cap X| \geqslant |B \cap X|$. Therefore (A, B) and (C, D) do not agree on X, which thus witnesses that (A, B) does not lie in the closure of $\tau \cap \vec{S}_k$ in S_k. $\square$

Note that in our proof above that (iii) implies (i) we did not make use of the assumption that the tangle τ was induced by an end of G: indeed, every orientation of S_k that has a finite decider set is closed in $\vec{S}_k$.

For an end tangle τ that is closed in $\vec{S}$ we can say slightly more about its decider sets in $\vec{S}_k$: for every $k \in \mathbb{N}$ the restriction $\tau \cap \vec{S}_k$ has a decider set of size exactly k which is a $(< k)$-inseparable set. Finding these $(< k)$-inseparable decider sets is straightforward: such an end tangle τ is defined by an $\aleph_0$-block K by Theorem 35, and every subset $X \subseteq K$ of size k is a $(< k)$-inseparable decider set for $\tau \cap \vec{S}_k$. However, having a $(< k)$-inseparable decider set for $\tau \cap \vec{S}_k$ for all $k \in \mathbb{N}$ is not a characterizing property for the closed end tangles of G:

Example 5.1.11. For $n \in \mathbb{N}$ let K^n be the complete graph on n vertices. Let G be the graph obtained from a ray $R = v_1 v_2 \dots$ by replacing each vertex v_n with the complete graph K^n, making each vertex from the K^n replacing v_n adjacent to all vertices from the K^{n+1} replacing v_{n+1}. Then G has a unique end ω; let τ be the end tangle induced by ω. Since ω is undominated τ is not closed in $\vec{S} = \vec{S}(G)$ by Theorem 36. However, for every $k \in \mathbb{N}$, the tangle $\tau \cap \vec{S}_k$ has a $(< k)$-inseparable decider set of size k: the clique K^k which replaced the vertex v_k of R is such a $(< k)$-inseparable decider set.

5.2 The tree-of-tangles theorem

5.2.1 Introduction

In Section 4.1 we established Theorem 18 and its canonical version Theorem 19, two unified theorems that imply virtually all previously known tree-of-tangles theorems for finite separation systems. The merit of Theorem 18 in particular lies in the fact that, while strong enough to imply all these results, the proof of the theorem is simple and its assumptions are easy-to-check.

Theorem 18 (Splinter theorem)**.** *Let U be a universe of separations and $\mathcal{A} = (A_i)_{i \leqslant n}$ a family of subsets of U. If $\mathcal{A}$ splinters then we can pick an element a_i from each A_i so that $\{a_1, \ldots, a_n\}$ is nested.*

Theorem 18, in a sense, is yet another step in a series of abstractions in the theory of tangles: rather than working with the tangles (or rather the more general profiles) themselves it operates just on the collection of sets of separations distinguishing a given pair of these.

Theorem 18 is proved by induction: it finds a separation $a_i \in \mathcal{A}_i$ which is nested with some element of each other $\mathcal{A}_j$, and then proceeds inductively on the remaining $n-1$ family members, restricted to those separations nested with a_i. This approach cannot deal with infinite families of sets, however.

In this section we overcome these difficulties and present two different ways to obtain a result akin to Theorem 18 for infinite families of sets of separations, each with its own set of assumptions. Both of these approaches use the concept of *compactness*: if a solution can be found for every finite sub-system, and these can be chosen compatibly locally, then there is a global solution for the infinite system as well.

To facilitate this approach we work with *profinite separation systems*, which are infinite separation systems whose structure is wholly determined by a family of finite separation systems. This concept generalises the structure of separations of infinite graphs: for these, too, one can recover all information from the knowledge of the separations of all finite subgraphs.

Our first approach, presented in Section 5.2.3, assumes that the sets $\mathcal{A}_i$ in Theorem 18 are closed in the profinite topology, and is otherwise unchanged:

Theorem 38. *Let $\vec{U} = \varprojlim (\vec{U}_p \mid p \in P)$ be a profinite universe of separations and $\mathcal{B}$ a family of non-empty closed subsets of $\vec{U}$. If $\mathcal{B}$ splinters then there is a nested set $N \subseteq \vec{U}$ containing at least one element from each member of $\mathcal{B}$.*

We apply this theorem in Section 5.2.4. The second approach for an infinite tree-of-tangles theorem assumes that the tangles themselves are closed sets:

Theorem 40. *Let $\vec{S} \subseteq \vec{U}$ be submodular and closed. Then there is a closed nested set $\vec{T} \subseteq \vec{S}$ that distinguishes all closed regular profiles of $\vec{S}$.*

We compare these two approaches in Section 5.2.5, and use our insights from that comparison to strengthen the application to graphs from Section 5.2.4 by combining the tools of both approaches.

5.2.2 Terminology and basic facts

In addition to the terms used in the study of separations so far, we will be using the following new or amended definitions and basic tools.

An *order-function* on a universe $\vec{U}$ is some function $|\cdot|\colon \vec{U} \to \mathbb{N}_0$ such that $|\vec{s}| = |\overleftarrow{s}| =: |s|$ for every unoriented separation $s \in U$. Note that, in contrast with the finite setting, we now require order-functions to take their values in $\mathbb{N}_0$ rather than in $\mathbb{R}_{\geqslant 0}$. In the finite case this difference is fairly immaterial: by the density of the rational numbers in the reals, any real-valued order function on a finite universe can be replaced by an, for all intents and purposes, equivalent integer-valued one. For finite universes this argument is no longer available, and we therefore need to ask explicitly that the order function be integer-valued.

Given a universe U with an order function we write S_k for the set of all separations in U with $|s| < k$, and we extend this notation to all $k \in \mathbb{N} \cup \{\aleph_0\}$. Of course, since $|\cdot|$ takes its values in $\mathbb{N}$, we have $\vec{S}_{\aleph_0} = \vec{U}$. However this notation will still be convenient in our applications: in a graph G we will use $\vec{S}_{\aleph_0}$ for the set of all separations (A, B) of G with finite $A \cap B$. (For graphs one could extend this notation to arbitrary cardinals, but we have no need to do so.)

In such a universe a *k-profile of U* is a profile of S_k in the usual sense, where we again allow $k \in \mathbb{N} \cup \{\aleph_0\}$. A *profile in U* is then a k-profile for some k. Two profiles in U are *distinguishable* if they are not subsets of each other. A profile P in U is *robust* if for all $\vec{r} \in P$ and $\vec{s} \in \vec{U}$ with $|\vec{r} \vee \vec{s}| < |r|$ and $|\vec{r} \vee \overleftarrow{s}| < |r|$ the profile P contains either $\vec{r} \vee \vec{s}$ or $\vec{r} \vee \overleftarrow{s}$.

We remark that in [17] robustness was defined in a slightly weaker but much more technical fashion, as a property of a set $\mathcal{P}$ of profiles in U. (For comparison, the definition of robustness given in Section 4.1.3 is a generalization of [17]'s.) However for the purpose of our upcoming proofs the difference between these two notions is insignificant. Every set $\mathcal{P}$ of pairwise distinguishable profiles is a *robust set of profiles* in the sense of [17] as soon as each profile in $\mathcal{P}$ is robust in our sense, and conversely a reader familiar with [17] may replace each appearance of a set of robust profiles here with a robust set of profiles without altering the proofs.

In our applications we will need the fact that the separations which efficiently distinguish a given pair of profiles exhibit a lattice-like structure:

Lemma 5.2.1. *Let $\vec{U}$ be a universe with a submodular order function and P and P' two profiles in $\vec{U}$. If $\vec{r}, \vec{s} \in P$ distinguish P and P' efficiently, then both $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ also lie in P and distinguish P and P' efficiently.*

Proof. If one of $\vec{r} \vee \vec{s}$ and $\vec{r} \wedge \vec{s}$ has order at most $|r| = |s|$, then that corner separation lies in P and distinguishes P and P' by their consistency and the property P. The efficiency of r and s now implies that neither of the two considered corner separations can have order strictly lower than $|r|$. Therefore, by submodularity, both of them have order exactly $|r|$, which implies the claim. □

In the remainder of this section we give the definitions and basic facts necessary for our applications in graphs. Let $G = (V, E)$ be a graph, which may (and generally will) be infinite.

Given a set $X \subseteq V$ of vertices we call a connected component C of $G - X$ *tight* if $N(C) = X$. For two vertices x and y of G, a set $X \subseteq V \setminus \{x, y\}$ with $|X| = k$ is an *x–y-separator of order k* if x and y lie in different components of $G - X$. We call $X \subseteq V \setminus \{x, y\}$ an *x–y-separator* if X is an x–y-separator or order $|X|$. Such a separator X is a *minimal* x–y-separator, or separates x and y *minimally*, if no proper subset of X is an x–y-separator. Observe that this minimality implies that the two components of $G - X$ containing x and y are tight. We will regularly use the following observation about these minimal separators:

Lemma 5.2.2 ([30, 2.4])**.** *Let G be a graph, $x, y \in V(G)$, and $k \in \mathbb{N}$. Then there are only finitely many separators of size at most k separating x and y minimally.*

5.2.3 The profinite splinter theorem

In this section we establish an extension of the Splinter Theorem to a large class of infinite separation systems: the profinite universes. Informally, a separation system is profinite if it is determined entirely by its finite subsystems. The most prominent, and most important, example of such a universe of separations is the separation system of an infinite graph: two separations of an infinite graph are comparable precisely if all their restrictions to finite subgraphs are comparable. Moreover, a pair (A, B) of sets of vertices of an infinite graph G is a separation if and only if the restriction of (A, B) to each finite subgraph H of G is a separation of H. We will make this relation between the separations of a profinite universe and their finite restrictions more formal in Section 5.2.3.

Introduction to profinite universes

For an in-depth introduction to profinite separation systems we refer the reader to [18], where this class of separation systems was first introduced. In this section we shall give only the definitions, terms, and tools for profinite universes relevant to our studies.

A *directed set* is a poset P in which every two elements have a common upper bound, i.e., in which there is an $r \in P$ with $p \leqslant r$ and $q \leqslant r$ for all $p, q \in P$. Given a directed set P, an *inverse system (of finite sets)* is a family $\mathcal{X} = (\, X_p \mid p \in P\,)$ of finite sets, together with maps $f_{qp}\colon X_q \to X_p$ for all $q > p$ that are *compatible* in the sense that $f_{rp} = f_{qp} \circ f_{rq}$ for all $r > q > p$. If every set X_p is a finite universe of separations $\vec{U}_p$, and the maps f_{qp} are homomorphisms of universes, then the family $\mathcal{U} = (\,\vec{U}_p \mid p \in P\,)$ is an *inverse system (of universes of separations).*

A *limit* of an inverse system $\mathcal{X} = (\, X_p \mid p \in P\,)$ is a compatible choice of one element x_p from each X_p, that is, a family $(\, x_p \mid p \in P\,)$ with $x_p \in X_p$ and $f_{qp}(x_q) = x_p$ for all $q > p$. The *inverse limit* $\varprojlim \mathcal{X}$ of $(\, X_p \mid p \in P\,)$ is the set of all limits of $\mathcal{X}$. It is a well-known fact that every inverse system of non-empty finite sets has a limit ([18]).

Limits and the *inverse limit* of an inverse system of universes are defined in the same way. Then the inverse limit $\vec{U} = \varprojlim \mathcal{U}$ of an inverse system of universes $\mathcal{U} = (\,\vec{U}_p \mid p \in P\,)$ is itself a universe of separations by defining involution, partial order, joins, and meets coordinate-wise. That is by, for $\vec{r} = (\,\vec{r}_p \mid p \in P\,)$ and $\vec{s} = (\,\vec{s}_p \mid p \in P\,)$, letting

$$\overleftarrow{s} := (\,\overleftarrow{s}_p \mid p \in P\,)$$

as well as

$$\vec{r} \vee \vec{s} := (\,\vec{r}_p \vee \vec{s}_p \mid p \in P\,)$$

and

$$\vec{r} \wedge \vec{s} := (\,\vec{r}_p \wedge \vec{s}_p \mid p \in P\,),$$

with $\vec{r} \leqslant \vec{s}$ if and only if $\vec{r}_p \leqslant \vec{s}_p$ for all $p \in P$. In particular the involution, joins and meets of limits of $\mathcal{U}$ are again limits of $\mathcal{U}$.

A universe of separations is then called *profinite* if it is isomorphic to the inverse limit of some finite universes of separations. The most prominent example of a profinite universe of separations comes from infinite graphs:

Example 5.2.3. Let $G = (V, E)$ be an infinite graph and $\vec{U} = \vec{U}(G)$ the universe of all separations of G. Then $\vec{U}$ is profinite: let $\mathcal{X}$ be the set of all finite $Z \subseteq V$. Then $\mathcal{X}$, ordered by inclusion, is a directed set. For $Z \in \mathcal{X}$ let $\vec{U}_Z$ be the universe of separations of $G[Z]$. We define maps $f_{ZY} \colon \vec{U}_Z \to \vec{U}_Y$ for $Y \subset Z$ by letting f_{ZY} map a separation (A_Z, B_Z) of $G[Z]$ to $(A_Z \cap Y\,,\, B_Z \cap Y)$, which is easily seen to be a separation of $G[Y]$. These maps are clearly compatible, and thus the family $\mathcal{U} \coloneqq (\,\vec{U}_Z \mid Z \in \mathcal{X}\,)$ is an inverse system of finite universes.

Let us show that $\vec{U}$ is isomorphic to the inverse limit of $\mathcal{U}$. For this observe that for every separation (A, B) of G the family of its restrictions $(\,(A \cap Z\,,\, B \cap Z) \mid Z \in \mathcal{X}\,)$ is a limit of $\mathcal{U}$, and the map $f \colon \vec{U} \to \varprojlim \mathcal{U}$ given by mapping each $(A, B) \in \vec{U}(G)$ to this family $(\,(A \cap Z\,,\, B \cap Z) \mid Z \in \mathcal{X}\,)$ is a homomorphism of universes. Moreover f is clearly injective, and its inverse is a homomorphism as well. To see that f is an isomorphism between $\vec{U}$ and $\varprojlim \mathcal{U}$ it thus remains to show that f is surjective, that is, that every limit of $\mathcal{U}$ gives rise to a separation of G.

So let $\vec{s} = (\,(A_Z, B_Z) \mid Z \in \mathcal{X}\,)$ be a limit of $\mathcal{U}$. Let A and B be the union of the sets A_Z and B_Z, respectively, over all $Z \in \mathcal{X}$. We claim that (A, B) is a separation of G. If so then $f((A, B)) = \vec{s}$, showing that f is surjective.

Note first that $A \cup B = V$, since $v \in A_{\{v\}} \cup B_{\{v\}}$ for every $v \in V$. Suppose now that G contains an edge $vw \in E$ with $v \in A \setminus B$ and $w \in B \setminus A$. Let $Z \coloneqq \{v, w\}$ and consider the induced subgraph $G[Z]$ of G: by definition of (A, B) we have $A_Z = \{v\}$ and $B_Z = \{w\}$, but then vw is an edge of $G[Z]$ between $A_Z \setminus B_Z$ and $B_Z \setminus A_Z$, contradicting the assumption that $(A_Z, B_Z) \in \vec{U}_Z$. Therefore (A, B) is indeed a separations of G.

For the remainder of this section, let $\mathcal{U} = (\,\vec{U}_p \mid p \in P\,)$ be an inverse system of universes and $\vec{U}$ its inverse limit.

Given an element $\vec{s} = (\,\vec{s}_p \mid p \in P\,)$ of $\vec{U}$, we write $(\vec{s} \restriction p) \coloneqq \vec{s}_p$ for the projection of $\vec{s}$ to $\vec{U}_p$. Likewise, for a set $O \subseteq \vec{U}$ the projection $O \restriction p$ to $\vec{U}_p$ is the set of all $\vec{s} \restriction p$ with $\vec{s} \in O$. We extend this notation to also include subsets of $\vec{U}_q$ for $q \in P$: for $q \geqslant p$ and $O_q \subseteq \vec{U}_q$ we also write $O_q \restriction p$ for $f_{qp}(O_q)$, where we understand f_{qp} to be the identity if $p = q$.

A family $(\,N_p \mid p \in P\,)$ of finite subsets of the $\vec{U}_p$ is a *restriction* of $\mathcal{U}$ if $f_{qp}(N_q) = N_p$ for all $q > p$. The inverse limit of such a restriction of $\mathcal{U}$ is a subset of $\vec{U} = \varprojlim \mathcal{U}$.

By equipping each $\vec{U}_p$ in $\mathcal{U}$ with the discrete topology, the inverse limit $\vec{U} = \varprojlim \mathcal{U}$ becomes a topological space as a subspace of the product space $\prod_{p \in P} \vec{U}_p$. Doing so makes the maps f_{qp} continuous, and it is easy to see that $\vec{U}$ is a closed subset of the product $\prod_{p \in P} \vec{U}_p$ and hence compact. In fact, the topology on $\vec{U}$ can be described in terms of the sets $\vec{U}_p$:

Lemma 5.2.4 ([18, Lemma 4.1])**.**
(i) *The topological closure in $\vec{U}$ of a set $O \subseteq \vec{U}$ is the set of all limits $\vec{s} = (\,\vec{s}_p \mid p \in P\,)$ with $\vec{s}_p \in O \restriction p$ for all p.*

(ii) *A set $O \subseteq \vec{U}$ is closed in $\vec{U}$ if and only if there are sets $O_p \subseteq \vec{U}_p$, with $f_{qp}(O_q) \subseteq O_p$ whenever $p < q \in P$, such that $O = \varprojlim (\, O_p \mid p \in P\,)$.*

We shall also use the following lemma, which is a re-formulation of [18, Lemma 5.4]:

Lemma 5.2.5. *A set $N \subseteq \vec{U}$ is nested if and only if $(N \restriction p) \subseteq \vec{U}_p$ is nested for all $p \in P$.*

A consequence of Lemma 5.2.4 and Lemma 5.2.5 (which we will not use) is that the topological closure of a nested set is still nested.

Finally, let us show that two closed sets in $\vec{U}$ intersect if all of their projections do:

Lemma 5.2.6. *If $A, B \subseteq \vec{U}$ are closed and $(A \restriction p) \cap (B \restriction p)$ is non-empty for every $p \in P$, then $A \cap B$ is non-empty.*

Proof. The family $(\,(A \restriction p) \cap (B \restriction p) \mid p \in P\,)$ is an inverse system of non-empty finite subsets of $\vec{U}$, using as maps the restrictions of the maps f_{qp} of $(\,\vec{U}_p \mid p \in P\,)$. Since both A and B are closed in $\vec{U}$, every limit of this family lies in $A \cap B$, which is therefore non-empty. $\square$

Stating and proving the profinite splinter theorem

Using the framework of profinite separation systems, we can now extend Theorem 18 to infinite separation systems:

Theorem 38. *Let $\vec{U} = \varprojlim (\,\vec{U}_p \mid p \in P\,)$ be a profinite universe of separations and $\mathcal{B}$ a family of non-empty closed subsets of $\vec{U}$. If $\mathcal{B}$ splinters then there is a nested set $N \subseteq \vec{U}$ containing at least one element from each member of $\mathcal{B}$.*

Proof. For $p \in P$ let $\mathcal{B}_p$ denote the projection $\mathcal{B} \restriction p$ of $\mathcal{B}$ to $\vec{U}_p$, that is, the family of all $B \restriction p$, where B is a member of $\mathcal{B}$. Then each projection $\mathcal{B}_p$ splinters in $\vec{U}_p$: consider two members B_p and B'_p of $\mathcal{B}_p$, with separations $\vec{r}_p \in B_p$ and $\vec{s}_p \in B'_p$. By definition of B_p and B'_p there are $\vec{r} \in B$ and $\vec{s} \in B'$ with $(\vec{r} \restriction p) = \vec{r}_p$ and $(\vec{s} \restriction p) = \vec{s}_p$. Since $\mathcal{B}$ splinters either at least one of $\vec{r}$ or $\vec{s}$ lies in $B \cap B'$, in which case its projection to $\vec{U}_p$ lies in $B_p \cap B'_p$, or else some corner separation $\vec{c}$ of $\vec{r}$ and $\vec{s}$ lies in $B \cup B'$. In the latter case $(\vec{c} \restriction p) \in B_p \cup B'_p$ is a corner separation of $\vec{r}_p$ and $\vec{s}_p$, showing that $\mathcal{B}_p$ indeed splinters.

By the above observation and Theorem 18 applied to $\mathcal{B}_p$ and $\vec{U}_p$ there is a nested set $N_p \subseteq \vec{U}_p$ for every $p \in P$ which contains an element of each member of $\mathcal{B}_p$. Let $\mathcal{N}_p$ be the set of all such nested sets $N_p \subseteq \vec{U}_p$. Observe that if $N_q \in \mathcal{N}_q$ and $q > p$ then $f_{qp}(N_q) \subseteq \vec{U}_p$ is a nested set meeting each member of $\mathcal{B}_p$ and hence lies in $\mathcal{N}_p$. Therefore the family $(\,\mathcal{N}_p \mid p \in P\,)$ together with the maps mapping $N_q \in \mathcal{N}_q$ to $f_{qp}(N_q) \in \mathcal{N}_p$ for $q > p$ is an inverse system of finite sets.

Let $(\,N_p \mid p \in P\,)$ be a limit of $(\,\mathcal{N}_p \mid p \in P\,)$. Then this limit is a restriction of $\mathcal{U}$, and hence $N := \varprojlim (\,N_p \mid p \in P\,)$ is a subset of $\vec{U}$. In fact N is a closed nested subset of $\vec{U}$ by Lemma 5.2.4 and Lemma 5.2.5. To see that N contains an element of each member of $\mathcal{B}$, let B be a member of $\mathcal{B}$. Then $(N \restriction p) \cap (B \restriction p)$ is non-empty for each $p \in P$ since $(N \restriction p) = N_p \in \mathcal{N}_p$, and thus by Lemma 5.2.6 and the assumption that B is closed in $\vec{U}$ the sets N and B intersect. $\square$

5.2.4 Applications of the profinite splinter theorem

For this section, let $G = (V, E)$ be a connected graph and $\mathcal{X}$ the set of finite subsets of V. As seen in Example 5.2.3, the universe $\vec{U} = \vec{U}(G)$ of separations of G is profinite since it arises as the inverse limit of $(\vec{U}_Z \mid Z \in \mathcal{X})$, where $\vec{U}_Z$ denotes the universe of separations of $G[Z]$. Following the notation of Section 5.2.3, we write $(A, B) \upharpoonright Z$ for the projection $(A \cap Z\,,\, B \cap Z)$ of a separation $(A, B) \in \vec{U}$ to $\vec{U}_Z$.

For $k \in \mathbb{N}$ let $\vec{S}_k = \vec{S}_k(G)$ be the separation system of all separations of order $< k$ of G. Using Lemma 5.2.4, it is easy to observe the following fact about these $\vec{S}_k$, which will be used throughout this section:

Observation 5.2.7. For every $k \in \mathbb{N}$ the set $\vec{S}_k$ is a closed subset of $\vec{U}$.

Our main goal in this section is to use Theorem 38 to find a nested set of separations which efficiently distinguishes a large set of profiles of G. Concretely, we will be able to distinguish the set of all *regular bounded* profiles in G. A profile P in G is *bounded* if P is a k-profile of G for some k but is not a subset of any $\aleph_0$-profile of G. Recall that a profile in G is regular if it contains no separation of the form $(V(G), X)$.

The main result of this section, then, will be the following:

Theorem 39. *Let $\mathcal{P}$ be a set of robust regular bounded profiles in G. Then there is a nested set N of separations of G which efficiently distinguishes all distinguishable profiles in $\mathcal{P}$.*

It can be shown ([11]) that every $\aleph_0$-profile P in G corresponds to either an end of G, or a so-called *ultrafilter tangle*: an orientation which, for some $X \in \mathcal{X}$, induces a non-principal ultrafilter on the set of components of $G - X$. These ultrafilter tangles are studied extensively in [11]. Ends and ultrafilter tangles both exhibit a different behaviour than profiles of finite graphs. Bounded profiles on the other hand *do* behave similarly to profiles of finite graphs, and consequently we shall be able to utilize Theorem 38 to establish Theorem 39.

By the above result of [11] every profile in G that is not bounded is the truncation of either an end tangle of G or an ultrafilter tangle. Conversely, if P is a maximal k-profile of G, i.e. if P does not extend to some $(k+1)$-profile of G, then P is clearly bounded. One way for this to happen is if P contains separations (A, B) and (C, D) with $V(G) = A \cup C$ and $|(A, B) \vee (C, D)| = k \geqslant 2$: in this case by the profile property any $(k+1)$-profile extending P would have to contain $(A, B) \vee (C, D)$, which is a co-small separation and hence cannot lie in any profile of order $\geqslant 3$ ([15]).

For the remainder of this section let $\mathcal{P}$ be a set of robust regular bounded profiles in G. Given two profiles P and P' in $\mathcal{P}$ let $\mathcal{A}_{P,P'}$ be the set of all separations of G that efficiently distinguish P and P'.

In order to deduce Theorem 39 from Theorem 38 we need to show that the family of the sets $\mathcal{A}_{P,P'}$ splinters, and that each $\mathcal{A}_{P,P'}$ is a closed subset of $\vec{U}$. We will start by showing the latter:

Proposition 5.2.8. *Let P and P' be distinguishable regular bounded profiles in G. Then $\mathcal{A}_{P,P'}$ is a closed subset of $\vec{U}$.*

In order to prove Proposition 5.2.8 we will first need to show a series of lemmas about how bounded profiles, and their efficient distinguishers, behave.

The first step is to show that a regular bounded profile, for every sufficiently small set X of vertices, points towards a component of $G - X$. That is to say: bounded profiles do not behave like the ultrafilter tangles of [11].

Lemma 5.2.9. *Let X be a finite set of vertices and P a regular bounded profile of order at least $|X| + 1$ in G. Then there is a unique component C of $G - X$ with $(V \setminus C\,,\, C \cup X) \in P$.*

Proof. Suppose that P contains for every component C of $G - X$ the separation $(C \cup X\,,\, V \setminus C)$, we are going to construct an extension of P to a profile of $S_{\aleph_0}$.

To determine the appropriate orientation of a separation $\{A, B\} \in S_{\aleph_0}$, consider the components of $G - X$ and let C_A be the union of all those components which are contained in $A \setminus B$. Likewise let C_B be the union of all components contained in $B \setminus A$ and C_R the union of the remaining components, i.e., those which meet both A and B. Since each of these needs to meet $A \cap B$ there are only finitely many.

By our assumption P contains for every component C of $G - X$ the separation $(C \cup X\,,\, V \setminus C)$. Since C_R is a union of only finitely many components of $G - X$ and since P is a profile, we have $(C_R \cup X\,,\, V \setminus C_R) \in P$.

We now want to prove that one of $(V \setminus C_B\,,\, C_B \cup N(C_B))$ and $(V \setminus C_A\,,\, C_A \cup N(C_B))$ lies in P. Indeed, if this is not the case then their respective inverses are in P, so the profile property gives us

$$(C_B \cup C_A \cup N(C_A \cup C_B)\,,\, V \setminus (C_A \cup C_B)) \in P.$$

This however would imply that the supremum

$$(C_B \cup C_A \cup N(C_A \cup C_B)\,,\, V \setminus (C_A \cup C_B)) \vee (C_R \cup X\,,\, V \setminus C_R) = (V, X)$$

lies in P by the profile property and the fact that this is a separation of order $|X|$. This contradicts the regularity of P.

This proves that one of $(V \setminus C_B\,,\, C_B \cup N(C_B))$ and $(V \setminus C_A\,,\, C_A \cup N(C_B))$ lies in P, and by consistency we cannot have both. We may thus define an orientation P' of $S_{\aleph_0}$ by declaring that (A, B) shall be in P' if and only if $(V \setminus C_B\,,\, C_B \cup N(C_B))$ is in P. This orientation is consistent since P is consistent and $(A, B) \leqslant (V \setminus C_B\,,\, C_B \cup N(C_B))$. Note that $P \subseteq P'$ and it only remains to show that P' is a profile.

Given $(A, B), (C, D) \in P'$ we have $(V \setminus C_B\,,\, C_B \cup N(C_B))$ and $(V \setminus C_D\,,\, C_D \cup N(C_D))$ in P, by definition. The profile property of P then gives us

$$(V \setminus C_B\,,\, C_B \cup N(C_B)) \vee (V \setminus C_D\,,\, C_D \cup N(C_D)) \in P \subseteq P',$$

so by the consistency of P' we have $((A, B) \vee (C, D))^* \notin P'$. □

Our next intermediate step is to show that bounded profiles cannot give rise to a decreasing sequence of components, each of which is the sole right side of a separation in that profile, and whose intersection is empty:

Lemma 5.2.10. *Let P be a k-profile of G and $(\,C_i \mid i \in I\,)$ an infinite sequence of non-empty connected vertex sets with $C_i \supseteq C_j$ for all $i \leqslant j$ and such that $(V \setminus C_i\,,\, C_i \cup N(C_i)) \in P$ for each $i \in I$. If $\bigcap_{i \in I} C_i$ is empty then P can be extended to an $\aleph_0$-profile of G.*

Proof. Let us define an orientation $\tilde{P}$ of $S_{\aleph_0}$ and show that $\tilde{P}$ is a profile extending P. Let $\tilde{P}$ consist of all $(A, B) \in \vec{S}_{\aleph_0}$ for which there is a C_i with $C_i \subseteq B \setminus A$. Since the C_i form a decreasing sequence of connected vertex sets with $\bigcap_{i \in \mathbb{N}} C_i = \emptyset$, and $A \cap B$ is finite, this orients each $\{A, B\} \in S_{\aleph_0}$. Moreover this orientation clearly is consistent.

For the profile property consider (A, B) and (A', B') in $\tilde{P}$. By definition there are C_i and C_j with $C_i \subseteq B \setminus A$ and $C_j \subseteq B' \setminus A'$. By symmetry we may assume that $i \leqslant j$ and thus $C_j \subseteq C_i$. But then $C_j \subseteq (B \setminus A) \cap (B' \setminus A')$, showing that $\tilde{P}$ contains $(A \cup A', B \cap B')$, as required.

To finish the proof it thus remains to verify that $P \subseteq \tilde{P}$. However any $(A, B) \in P$ with $C_i \subseteq A \setminus B$ for some C_i would be inconsistent with $(V \setminus C_i , C_i \cup N(C_i)) \in P$. Therefore $\tilde{P}$ is an extension of P. □

Using Lemma 5.2.9 and 5.2.10 we can take the next step towards showing that the sets $\mathcal{A}_{P,P'}$ are closed by showing that for every infinite chain in $\mathcal{A}_{P,P'}$ the sequence of the separators is eventually constant:

Lemma 5.2.11. *Let P and P' be distinguishable bounded profiles in G and $(A_1, B_1) \leqslant (A_2, B_2) \leqslant \ldots$ an infinite increasing sequence in $\mathcal{A}_{P,P'}$. Then the sequence $(A_i \cap B_i)_{i \in \mathbb{N}}$ is eventually constant.*

Proof. By switching their roles if necessary we may assume that P' contains (B_1, A_1). Then, by consistency, $(B_i, A_i) \in P'$ and consequently $(A_i, B_i) \in P$ for every $i \in \mathbb{N}$.

For $i \in \mathbb{N}$ let us write $X_i := A_i \cap B_i$. By Lemma 5.2.9 there is a unique component C_i of $G - X_i$ with $(V \setminus C_i , C_i \cup X_i) \in P$. Observe that, just like (A_i, B_i), the separation $(V \setminus C_i , C_i \cup X_i)$ distinguishes P and P' efficiently. This efficiency implies that $N(C_i) = X_i$. Observe further that the separations $(V \setminus C_i , C_i \cup X_i)$ form an increasing chain in $\mathcal{A}_{P,P'}$ whose sequence of separators is $(A_i \cap B_i)_{i \geqslant \mathbb{N}}$. In fact $C_i \supseteq C_j$ for $i \leqslant j$. It thus suffices to show that the sequence of the C_i is eventually constant. So suppose for a contradiction that the sequence of the C_i is strictly decreasing.

To obtain a contradiction it suffices by Lemma 5.2.10 to show that $\bigcap_{i \in \mathbb{N}} C_i$ is empty. So suppose that there is a vertex $v \in \bigcap_{i \in \mathbb{N}} C_i$. By applying Lemma 5.2.9 to P' and each X_i, we obtain components C'_i of $G - X_i$ such that $(V \setminus C'_i , C'_i \cup X_i) \in P'$ for all $i \in \mathbb{N}$. Clearly $C_i \neq C'_i$ and $N(C'_i) = X_i$ for all $i \in \mathbb{N}$. Fix any $w \in C'_1$. Then $w \in C'_i$ for every $i \in \mathbb{N}$, and each X_i is a minimal v-w-separator in G. The latter contradicts Lemma 5.2.2 by the assumption that all C_i, and hence all X_i, are distinct. Thus $\bigcap_{i \in \mathbb{N}} C_i$ is indeed empty, from which we can derive a contradiction to the boundedness of P using Lemma 5.2.10. □

Moreover, we can even show that the same statement also holds not only for chains of separations, but even for the entire set $\mathcal{A}_{P,P'}$:

Lemma 5.2.12. *Let P and P' be distinguishable regular bounded profiles in G. Then the separations $(A, B) \in \mathcal{A}_{P,P'}$ have only finitely many distinct separators $A \cap B$.*

Proof. Suppose for a contradiction that $\mathcal{A}_{P,P'}$ contains an infinite sequence of separations $(A_1, B_1), (A_2, B_2), \ldots$ whose separators $A_i \cap B_i$ are pairwise distinct. We may assume without loss of generality that $(A_i, B_i) \in P$ for every $i \in \mathbb{N}$. By Lemma 5.2.1 P contains all finite joins of these separations. For each $i \in \mathbb{N}$

let X_i be the separator of the supremum of (A_1, B_1) up to (A_i, B_i), and let C_i be the component of $G - X_i$ with $(V \setminus C_i\,, C_i \cup X_i) \in P$ as given by Lemma 5.2.9. By Lemma 5.2.11 the sequence of the X_i is eventually constant, and therefore so is the sequence of the C_i as P is a profile. Let $C := \bigcap_{i \in \mathbb{N}} C_i \neq \emptyset$.

Analogously for P' let X'_i be the separator of the supremum of (B_1, A_1) up to (B_i, A_i) and let C'_i be the component of $G - X'_i$ with $(V \setminus C'_i\,, C'_i \cup X'_i) \in P'$. As before let $C' := \bigcap_{i \in \mathbb{N}} C'_i \neq \emptyset$.

Since C_1 and C'_1 are disjoint so are C and C'. Fix vertices $v \in C$ and $w \in C'$. We claim that every separator $A_i \cap B_i$ is a minimal v-w-separator in G, contradicting the assertion of Lemma 5.2.2. To see this, consider $A_i \cap B_i$ for some $i \in \mathbb{N}$. Let $\tilde{C}_i$ and $\tilde{C}'_i$ be the components of $G - (A_i \cap B_i)$ obtained by applying Lemma 5.2.9 with P and P', respectively. Then $C \subseteq C_i \subseteq \tilde{C}_i$ and likewise $C' \subseteq C'_i \subseteq \tilde{C}'_i$, giving $v \in \tilde{C}_i$ and $w \in \tilde{C}'_i$. Moreover both $\tilde{C}_i$ and $\tilde{C}'_i$ have all of $A_i \cap B_i$ as their neighbourhood as (A_i, B_i) efficiently distinguishes P and P', and hence $A_i \cap B_i$ is indeed a minimal v-w-separator in G. □

We are now ready to prove Proposition 5.2.8, i.e. that the sets $\mathcal{A}_{P,P'}$ are closed subsets of $\vec{U}$:

Proof of Proposition 5.2.8. Let two distinguishable regular bounded profiles P and P' in G be given. By Lemma 5.2.12 only finitely many sets, say $X_1, \dots, X_n$, appear as separators of separations in $\mathcal{A}_{P,P'}$. For each X_i let C_i and C'_i be the two components of $G - X_i$ given by applying Lemma 5.2.9 to X_i for P and P', respectively.

We are now able to give a complete description of the set $\mathcal{A}_{P,P'}$: it is easy to check that a separation $(A, B) \in \vec{S}_{\aleph_0}$ distinguishes P and P' efficiently if and only if $A \cap B = X_i$ for some i with one of C_i and C'_i being a subset of A and the other a subset of B.

For each X_i, the set of all $(A, B) \in \vec{U}$ with $A \cap B = X_i$ as well as $C_i \subseteq A$ and $C'_i \subseteq B$ is closed by Lemma 5.2.4. Likewise the set of all $(A, B) \in \vec{U}$ with separator X_i as well as $C'_i \subseteq A$ and $C_i \subseteq B$ is closed, too. Therefore $\mathcal{A}_{P,P'}$ is the union of finitely many closed subsets of $\vec{U}$ and hence closed. □

Having established that the sets $\mathcal{A}_{P,P'}$ are closed in $\vec{U}$, it thus remains for us to verify that the family of the sets $\mathcal{A}_{P,P'}$ splinters in order to deduce Theorem 39 from Theorem 38. Since we shall need a slightly stronger property than splintering at a later point in Section 5.2.5, we will prove this stronger assertion here. In particular we will not make use of the assumptions that the profiles in $\mathcal{P}$ are regular and bounded.

To show that the sets $\mathcal{A}_{P,P'}$ splinter, we need to show that for all $(A, B) \in \mathcal{A}_{P,P'}$ and $(C, D) \in \mathcal{A}_{Q,Q'}$, either some corner separation of (A, B) and (C, D) lies in $\mathcal{A}_{P,P'}$ or in $\mathcal{A}_{Q,Q'}$, or one of (A, B) and (C, D) lies in both $\mathcal{A}_{P,P'}$ and $\mathcal{A}_{Q,Q'}$. In fact we will show that the first option always occurs.

We will split our proof of this into two separate lemmas, distinguishing the cases of equal and of distinct order of (A, B) and (C, D).

Let us first deal with the case that (A, B) is of strictly lower order than (C, D). In this case we can say precisely which of $\mathcal{A}_{P,P'}$ and $\mathcal{A}_{Q,Q'}$ will contain a corner separation of (A, B) and (C, D):

Lemma 5.2.13. *Let $(A, B) \in \mathcal{A}_{P,P'}$ and $(C, D) \in \mathcal{A}_{Q,Q'}$ with $|(A, B)| < |(C, D)|$. Then some corner separation of (A, B) and (C, D) lies in $\mathcal{A}_{Q,Q'}$.*

Proof. Since $|(A,B)| < |(C,D)|$ it follows that both Q and Q' orient $\{A,B\}$ the same, say $(A,B) \in Q \cap Q'$. If $|(A,B) \vee (C,D)| \leqslant |(C,D)|$ or $|(A,B) \vee (D,C)| \leqslant |(C,D)|$, it follows that this corner separation efficiently distinguishes Q and Q' by Lemma 2.1.1, so suppose that this is not the case. Then submodularity implies that $|(B,A) \vee (C,D)| < |(A,B)|$ and $|(B,A) \vee (D,C)| < |(A,B)|$, which in turn contradicts the efficiency of (A,B), since one of $(B,A) \vee (C,D)$ and $(B,A) \vee (D,C)$ would also distinguish the two robust profiles P and P'. □

For separations r and s the corner separations given by $\vec{r} \vee \vec{s}$ and $\overleftarrow{r} \vee \overleftarrow{s}$ (as well as their underlying unoriented separations) are referred to as *opposite* corner separations.

The second case is that (A,B) and (C,D) are of equal order. Here we can show that there are two opposite corner separations of (A,B) and (C,D) that lie in $\mathcal{A}_{P,P'}$ or in $\mathcal{A}_{Q,Q'}$:

Lemma 5.2.14. *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 5.2.1 gives $(A,B) \vee (C,D) \in \mathcal{A}_{P,P'}$ and $(B,A) \vee (D,C) \in \mathcal{A}_{P,P'}$. Thus by property P we also have $(A,B) \vee (C,D) \in \mathcal{A}_{Q,Q'}$, producing the desired pair of opposite corner separations. If Q' contains (B,A) we argue analogously.

So suppose that $(C,D) \in P'$ and $(A,B) \in Q'$. Then $(B,A) \vee (C,D) \in P'$ and $(A,B) \vee (D,C) \in Q'$ by the profile property, since by submodularity and the efficiency of (A,B) and (C,D) both of these corner separations have order exactly $|(A,B)|$. These two separations, then, are opposite corner separations of (A,B) and (C,D) with the first lying in $\mathcal{A}_{P,P'}$ and the second lying in $\mathcal{A}_{Q,Q'}$.

The remaining case is that no two of the four profiles agree in their orientation of $\{A,B\}$ and $\{C,D\}$. But then both of (A,B) and (C,D) lie in $\mathcal{A}_{P,P'}$ as well as in $\mathcal{A}_{Q,Q'}$, and the existence of two pairs of opposite corner separations, one with both elements in $\mathcal{A}_{P,P'}$ and one with both in $\mathcal{A}_{Q,Q'}$, follows from Lemma 5.2.1 and the disagreement of the four profiles on $\{A,B\}$ and $\{C,D\}$. □

We now have all the ingredients necessary for a proof of Theorem 39:

Proof of Theorem 39. By Proposition 5.2.8, we can apply Theorem 38. Thus we only need to show that the collection of these sets $\mathcal{A}_{P,P'}$ splinters. However, this follows from Lemma 5.2.13 and Lemma 5.2.14. □

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:

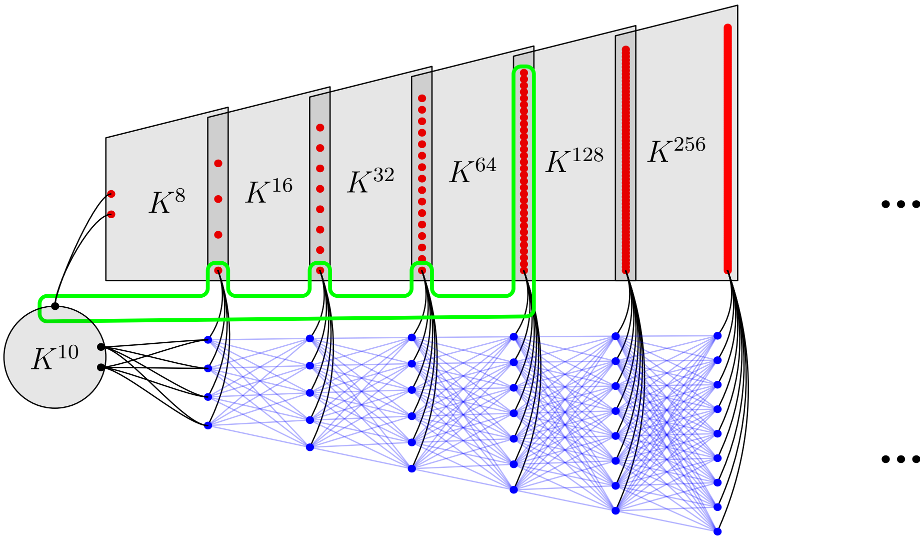

Figure 5.1: A locally finite graph where no tree-decomposition distinguishes all the finite-order tangles efficiently. The green separation is the only separation efficiently distinguishing the tangle induced by the K^{64} and the tangle induced by the K^{128}.

Example 5.2.15. Consider the graph displayed in Fig. 5.1. 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 $\frac{2}{3}\cdot 2^{n+1}$. The only separation which efficiently distinguishes P_n and P_{n+1} is the separation s_n with separator $\{v_1^i \mid i < n\} \cup \{u_j^{n+1}\}$.

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 bounded profiles efficiently: the separations given by such a tree-decomposition would have to contain this chain of order type $\omega + 1$, and it was shown in [29] that a nested set of nontrivial separations corresponds to a tree-decomposition if and only if it does not contain a chain of order type $\omega + 1$.

5.2.5 A tree-of-tangles theorem for ends

We conclude this chapter by giving another infinite tree-of-tangles theorem based on the concept of profinite separation systems. However, our approach to this will differ from that of Theorem 38: rather than asking that the sets of distinguishing separations of the profiles be closed we shall demand that the profiles themselves be closed in the profinite topology. We will thus apply the principle of compactness to these profiles directly rather than to their symmetric differences.

For the remainder of this section let $\vec{U} = \varprojlim(\vec{U}_p \mid p \in P)$ be a profinite universe with directed set P and bonding maps $f_{qp}\colon \vec{U}_q \to \vec{U}_p$ for $q > p$. The result we establish is the following extension of Theorem 16:

Theorem 40. *Let $\vec{S} \subseteq \vec{U}$ be submodular and closed. Then there is a closed nested set $\vec{T} \subseteq \vec{S}$ that distinguishes all closed regular profiles of $\vec{S}$.*

Preliminary results

For simplicity we shall assume without loss of generality that the inverse system $(\vec{U}_p \mid p \in P)$ is surjective, that is, that each bonding map f_{qp} is surjective. For the remainder of this section let $\vec{S} \subseteq \vec{U}$ be a separation system that is closed in $\vec{U}$. For $p \in P$ we shall write $\vec{S}_p$ for the projection $\vec{S} \restriction p$ of $\vec{S}$ to $\vec{U}_p$. Observe that $\vec{S}$ is submodular (in $\vec{U}$) if and only of for all $p \in P$ the projection $\vec{S}_p$ is submodular in $\vec{U}_p$.

If $\vec{S}$ contains a degenerate separation then there are no profiles of S; since we are interested in the profiles of S we shall therefore assume from now on that no element of $\vec{S}$ is degenerate. We can then make without loss of generality the further assumption that no projection $\vec{S}_p$ of $\vec{S}$ contains a degenerate separation either: for if, again and again, the $\vec{S}_p$ would contain a degenerate separation, then so would their inverse limit $\vec{S}$.

Our first lemma says that an orientation of S is closed in $\vec{S}$ (equivalently: in $\vec{U}$) precisely if for sufficiently large p its projection to $\vec{S}_p$ is an orientation of S_p:

Lemma 5.2.16. *Let O be an orientation of S. Then the following are equivalent:*

(i) *O is closed in $\vec{S}$;*

(ii) *$O \restriction p$ is an orientation of S_p for some $p \in P$;*

(iii) *$O \restriction p$ is an orientation of S_p for every $p \geqslant p_0$ for some $p_0 \in P$.*

Proof. $(iii) \implies (ii)$ is trivial.

For $(ii) \implies (i)$ suppose that O is not closed in $\vec{S}$. Then there is some $\vec{s} \in O$ for which $\overleftarrow{s}$ lies in the closure of O in $\vec{S}$. But then $O \restriction p$ contains both $\vec{s} \restriction p$ and $\overleftarrow{s} \restriction p$ for each $p \in P$, showing that $O \restriction p$ is never an orientation of S_p.

For $(i) \implies (iii)$ suppose that (iii) fails and consider for each $p \in P$ the subset of $\vec{S}_p$ of those separations that hove both of their orientations in $O \restriction p$. By assumptions these sets are non-empty, and hence from an inverse system of sets whose limit points lie in O and witness that O is not antisymmetric. $\square$

Let us now assume that $\vec{S}$ is submodular. Our next lemma asserts that consistency, like antisymmetry above, can be passed on from a closed orientation of S to its projections:

Lemma 5.2.17. *Let O be a consistent orientation of S, and $p \in P$ such that $O \restriction p$ is an orientation of S_p. If O is regular then $O \restriction p$ is consistent.*

Moreover if $S = U$ the assumption that O be regular can be dropped.

Proof. Let us show the main assertion. For this suppose that $O \restriction p$ is inconsistent. Then there are $\vec{r}$ and $\vec{s}$ in O such that $\{\vec{r}, \vec{s}\} \restriction p$ is an inconsistent pair in $\vec{S}_p$, that is, such that $(\overleftarrow{r} \restriction p) \leqslant (\vec{s} \restriction p)$ with $(r \restriction p) \neq (s \restriction p)$.

Since $\vec{S}$ is submodular one of $\overleftarrow{r} \wedge \vec{s}$ and $\vec{r} \wedge \overleftarrow{s}$ lies in $\vec{S}$; by symmetry we may assume that $(\overleftarrow{r} \wedge \vec{s}) \in \vec{S}$. Then $(\overleftarrow{r} \wedge \vec{s}) \leqslant \vec{s}$ and hence $(\overleftarrow{r} \wedge \vec{s}) \in O$ by consistency unless $(\overleftarrow{r} \wedge \vec{s}) = \overleftarrow{s}$. However the latter case is impossible since then $\overleftarrow{s}$ would be small by $\overleftarrow{s} = (\overleftarrow{r} \wedge \vec{s}) \leqslant \vec{s}$, contradicting the regularity of O. Therefore $(\overleftarrow{r} \wedge \vec{s}) \in O$.

Using that $(\overleftarrow{r} \restriction p) \leqslant (\vec{s} \restriction p)$ we now find

$$(\overleftarrow{r} \wedge \vec{s}) \restriction p = (\overleftarrow{r} \restriction p) \wedge (\vec{s} \restriction p) = \overleftarrow{r} \restriction p\,,$$

which shows that $O \restriction p$ is not antisymmetric, contrary to assumption.

Consider now the above arguments in the special case that $S = U$. In this case both of $(\overleftarrow{r} \wedge \vec{s})$ and $(\vec{r} \wedge \overleftarrow{s})$ lie in $\vec{S}$. If we can apply consistency to show that one of them lies in O then the argument is the same as above; otherwise $(\overleftarrow{r} \wedge \vec{s}) = \overleftarrow{s}$ as well as $(\vec{r} \wedge \overleftarrow{s}) = \overleftarrow{r}$. But if both those equalities hold then $\overleftarrow{s} \leqslant \overleftarrow{r}$ and $\overleftarrow{r} \leqslant \overleftarrow{s}$, so $\overleftarrow{r} = \overleftarrow{s}$, contradicting that their projections are an inconsistent pair. Therefore the statement holds for $\vec{S} = \vec{U}$ even without the assumption that O contains all small separations. $\square$

Recall that a subset $\mathcal{P}$ of $\vec{S}$ has the profile property P if there are no $\vec{r}, \vec{s}, \vec{t} \in \mathcal{P}$ with $\vec{r} \vee \vec{s} = \overleftarrow{t}$. If $\mathcal{P}$ is closed then this property, too, is inherited by sufficiently large projections:

Lemma 5.2.18. *Let $\mathcal{P} \subseteq \vec{S}$ be a closed subset with the profile property. Then for some $p_0 \in P$ each projection $\mathcal{P} \restriction p$ with $p \geqslant p_0$ has the profile property.*

Proof. Suppose not. Then for every $p \in P$ there are $q \geqslant p$ and $\vec{r}, \vec{s}, \vec{t} \in (\mathcal{P} \restriction q)$ with $\vec{r} \vee \vec{s} = \overleftarrow{t}$. Note that for such a triple we also have $(\vec{r} \restriction p) \vee (\vec{s} \restriction p) = \overleftarrow{t} \restriction p$. Let B_p be the set of all triples $(\vec{r}, \vec{s}, \vec{t})$ such that $\vec{r}, \vec{s}, \vec{t} \in \mathcal{P} \restriction p$ with $\vec{r} \vee \vec{s} = \overleftarrow{t}$. These form an inverse system of non-empty finite sets. Since $\mathcal{P}$ is closed in $\vec{S}$ it must contain the elements of a compatible choice in this system, which shows that $\mathcal{P}$ does not have the profile property. $\square$

Essentially, all three of these lemmas 5.2.16, 5.2.17, and 5.2.18 show that finitary properties of closed subsets of $\vec{S}$ pass onto sufficiently large projections: if a counterexample to some property consists of a fixed number of separations, and projections of counterexamples are again counterexamples, then a closed subset of $\vec{S}$ has that property if and only if each projection to $\vec{S}_p$ does for $p \geqslant p_0$ with some $p_0 \in P$. Antisymmetry, consistency, and the profile property are all examples of such finitary properties.

We can combine Lemma 5.2.16, 5.2.17 and 5.2.18 into the following:

Lemma 5.2.19. *Let $\mathcal{P}$ be a closed profile of $\vec{S}$. If $\mathcal{P}$ is regular, or else if $S = U$, then there is a $p_0 \in P$ such that $\mathcal{P} \restriction p$ is a profile of S_p for every $p \geqslant p_0$.* $\square$

Proving the tree-of-tangles theorem

We are now ready to establish the main result of this section:

Theorem 40. *Let $\vec{S} \subseteq \vec{U}$ be submodular and closed. Then there is a closed nested set $\vec{T} \subseteq \vec{S}$ that distinguishes all closed regular profiles of $\vec{S}$.*

Proof. As in the previous section we write $\vec{S}_p$ for the projection $\vec{S} \restriction p$ of $\vec{S}$ to $\vec{U}_p$.

If some element of S is degenerate then S has no profiles and the statement holds. Therefore we may assume that S contains no degenerate separation. Since $\vec{S}$ is closed we may further assume without loss of generality that no $\vec{S}_p$ has a degenerate element.

For each $p \in P$ let $\mathcal{T}_p$ consist of all sets of the form $\vec{T}_q \restriction p$, where $q \geqslant p$ and $\vec{T}_q$ is a nested subset of $\vec{S}_q$ that distinguishes all profiles of S_q. Then $(\,\mathcal{T}_p \mid p \in P\,)$, with element-wise bonding maps inherited from $\vec{U}$, is an inverse system of finite sets. From Lemma 5.2.19 and Theorem 16 it follows that $\mathcal{T}_p$ is non-empty for each $p \in P$.

Let $\vec{T} = (\,\vec{T}_p \mid p \in P\,)$ be a compatible choice in this system. Then we can view $\vec{T}$ as a subset of $\vec{U}$. In fact, by construction, we have that $\vec{T}$ is a closed subset of $\vec{S}$. Furthermore Lemma 5.2.5 tells us that $\vec{T}$ is nested since each $\vec{T}_p$ is. It remains to check that $\vec{T}$ distinguishes all closed regular profiles of S.

To see this let $\mathcal{P}$ and $\mathcal{P}'$ be two closed regular profiles of S. Pick $p_0 \in P$ large enough that $\mathcal{P} \restriction p$ and $\mathcal{P}' \restriction p$ are distinct profiles of S_p for every $p \geqslant p_0$; this is possible by Lemma 5.2.19. Pick any $p \geqslant p_0$. By definition there are $q \geqslant p$ and a nested set $\vec{T}'_q \subseteq \vec{S}_q$ that $\vec{T}'_q$ distinguishes all profiles of $\vec{S}_q$ such that $\vec{T}'_q \restriction p = \vec{T}_p$. Since $\mathcal{P} \restriction q$ and $\mathcal{P}' \restriction q$ are distinct profiles of S_q there is some $\vec{t}_q \in \vec{T}'_q$ with $\vec{t}_q \in (\mathcal{P} \restriction q)$ and $\overleftarrow{t}_q \in (\mathcal{P}' \restriction q)$. Then also $(\vec{t}_q \restriction p) \in (\mathcal{P} \restriction p)$ and $(\overleftarrow{t}_q \restriction p) \in (\mathcal{P}' \restriction p)$, with both of $\vec{t}_q \restriction p$ and $\overleftarrow{t}_q \restriction p$ lying in $\vec{T}_p = \vec{T}'_q \restriction p$. Therefore $\vec{T}$ contains some t with $t \restriction p = t_q \restriction p$. This separation t then distinguishes $\mathcal{P}$ and $\mathcal{P}'$. $\square$

If one so desires one can delete all non-regular separations from T to obtain a regular tree set that still distinguishes all closed regular profiles of S. However this tree set will in general no longer be closed. Indeed, any closed set containing an infinite star also contains a small separation (cf. [18]).

Theorem 40 adapts in the special case of $S = U$ as follows, with the same proof:

Theorem 41. *Let $\vec{U}$ be a profinite universe. Then there is a closed nested subset $\vec{T} \subseteq \vec{U}$ that distinguishes all closed profiles of U.*

Closed and bounded ends and tangles

Despite their difference in approach, it turns out that Theorem 40 can, after all, be derived from Theorem 38. In addition to obtaining another proof of Theorem 40, the tools we develop for this purpose will also enable us to strengthen the application Theorem 39 of Theorem 38 to apply to all *closed* profiles in a graph rather than to all *bounded* ones. It is nontrivial to see that this is in fact a strengthening, i.e. that all bounded profiles are closed, and we shall prove this assertion at the end of this section.

Let $\vec{U} = \varprojlim\,(\vec{U}_p \mid p \in P)$ be a profinite universe with directed set P and bonding maps $f_{qp} \colon \vec{U}_q \to \vec{U}_p$ for $q > p$. To deduce Theorem 40 from Theorem 38 we need to show that if two profiles Q and Q' of a closed $S \subseteq U$ are closed, then the set $A_{Q,Q'} \subseteq \vec{U}$ of all elements of S that distinguish the two is closed as well. Since Q and Q' are orientations of the same set S, we have that $A_{Q,Q'}$ is precisely the symmetric difference $Q \Delta Q'$ of Q and Q'. It is not true in general that the symmetric difference of closed sets is closed; however, in our case, it is true.

Let us establish the slightly stronger statement that the separations distinguishing any two closed antisymmetric sets form a closed set:

Proposition 5.2.20. *Let O and O' be closed antisymmetric subsets of $\vec{U}$. Then the set $A_{O,O'}$ of all $\vec{s} \in \vec{U}$ with $\vec{s} \in O$ and $\overleftarrow{s} \in O'$ or vice-versa is closed in $\vec{U}$.*

Proof. It suffices to show that $O \cap A_{O,O'}$ is closed since $A_{O,O'}$ is the union of $O \cap A_{O,O'}$ and $O' \cap A_{O,O'}$. Clearly $O \cap A_{O,O'}$ is the set of all $\vec{s} \in O$ with $\overleftarrow{s} \in O'$. By Lemma 5.2.4 we need to show that if $\vec{r} \in \vec{U}$ is such that $(\vec{r} \restriction p) \in (O \cap A_{O,O'}) \restriction p$ for each $p \in P$ then $\vec{r} \in O$ and $\overleftarrow{r} \in O'$.

For this observe that if $(\vec{r} \restriction p) \in (O \cap A_{O,O'}) \restriction p$ for some $p \in P$ then $(\vec{r} \restriction p) = (\vec{s} \restriction p)$ for some $\vec{s} \in O$ with $\overleftarrow{s} \in O'$, and consequently $(\overleftarrow{r} \restriction p) = (\overleftarrow{s} \restriction p) \in O' \restriction p$. Therefore, if $(\vec{r} \restriction p) \in (O \cap A_{O,O'}) \restriction p$ for all $p \in P$, then $(\vec{r} \restriction p) \in O \restriction p$ and $(\overleftarrow{r} \restriction p) \in O' \restriction p$ for all $p \in P$ as well, giving $\vec{r} \in O$ and $\overleftarrow{r} \in O'$ since O and O' are closed. We therefore have $\vec{r} \in O \cap A_{O,O'}$, showing that this set is closed. $\square$

Note that the set $A_{O,O'}$ defined in the assertion of Proposition 5.2.20 is not quite the same as the set of all $\vec{s} \in \vec{U}$ that distinguish O and O': it might contain degenerate separations from $O \cap O'$. This makes no difference in practice, however, since profiles cannot contain degenerate separations.

Using Proposition 5.2.20 and the above observation we can now obtain Theorem 40 from Theorem 38 by applying the latter to the family $\mathcal{A} = (A_{Q,Q'} \mid Q \neq Q')$, where the Q and Q' range over all closed regular profiles of $S \subseteq U$:

Second proof of Theorem 40. For each pair Q, Q' of closed regular profiles of $S \subseteq U$ let $A_{Q,Q'}$ be defined as in Proposition 5.2.20; these sets are closed. Since profiles do not contain degenerate separations the set $A_{Q,Q'}$ is precisely the set of separations in $\vec{S}$ which distinguish Q and Q'. The family $\mathcal{A} = (A_{Q,Q'} \mid Q \neq Q')$ splinters (cf. for instance Lemma 4.1.2), and the claimed nested set is now provided by Theorem 38. $\square$

Let us now utilize Proposition 5.2.20 to strengthen Theorem 39. For the remainder of this section let $G = (V, E)$ be an infinite graph. Recall that a profile in G is a k-profile of G for some integer k. If P is a k-profile of G then P is closed in $\vec{U} = \vec{U}(G)$ precisely if P is a closed subset of $\vec{S}_k \subseteq \vec{U}$, where $\vec{U}$ is equipped with the profinite topology; this uses the fact that each $\vec{S}_k$ is a closed subset of $\vec{U}$. In Section 5.2.4 we showed the following for G:

Theorem 39. *Let $\mathcal{P}$ be a set of robust regular bounded profiles in G. Then there is a nested set N of separations of G which efficiently distinguishes all distinguishable profiles in $\mathcal{P}$.*

With this new tool becoming available we can now upgrade to read 'closed' instead of 'bounded':

Theorem 42. *Let $\mathcal{P}$ be a set of robust regular profiles of G that are closed in $\vec{S}_{\aleph_0}(G)$. Then there is a nested set N of separations of G which efficiently distinguishes all distinguishable profiles in $\mathcal{P}$.*

Apart from changing 'bounded' to 'closed' – which we shall discuss in a moment – this Theorem 42 contains another subtle strengthening: it includes tangles of $\vec{S}_{\aleph_0}$, provided they are closed in $\vec{S}_{\aleph_0}$ with the subspace topology of $\vec{U}(G)$. Note that $\vec{S}_{\aleph_0}$ is never closed in $\vec{U}(G)$ if G is infinite, since (V, V)

always lies in its closure in $\vec{U}$. Therefore $\vec{S}_{\aleph_0}$ is not itself profinite. Theorem 42 is nevertheless able to handle profiles of infinite order which are closed in $\vec{S}_{\aleph_0}$ though: the set $\mathcal{A}_{Q,Q'}$ of separations efficiently distinguishing the profiles $Q, Q' \in \mathcal{P}$ is always contained in some $\vec{S}_k$, which is closed in $\vec{U}(G)$, and the restrictions of Q and Q' to this $\vec{S}_k$ are then closed as well.

Both of Theorem 39 and Theorem 42 are actually included in the stronger result [4, Theorem 5.12]. There the assumption the the profiles be bounded or closed is replaced by the weaker one that they obey the assertion of Lemma 5.2.9: that for each set X of fewer than k vertices the k-profile points to some component of $G - X$. Such profiles are sometimes called *principal* in the literature. Readers wishing to study [4] should be alert to the fact that [4] uses a non-standard definition of profile: namely, all profiles considered in that work are assumed to be principal in the above sense. [4, Theorem 5.12]'s inclusion of bounded profiles is immediate from Lemma 5.2.9, and it is not difficult to see that closed profiles are principal as well. It is similarly straightforward to check that [4, Theorem 5.12], which as stated only includes k-profiles for integer k, extends to principal profiles of $\vec{S}_{\aleph_0}$ just as Theorem 42 does.

Our short derivations of Theorem 39 and Theorem 42 still hold value, though, since they employ the general-purpose theorem 38 and are fairly compact, whereas [4] is a large body of work and treats the special case of separations of graphs only. We remark that due to the specialisation in graphs [4, Theorem 5.12] is out of reach for any approach relying on Theorem 38: if G is the double ray and τ_1, τ_2 are the end tangles of its two ends, then the set A_{τ_1,τ_2} of all separations which efficiently distinguish them is not closed in $\vec{S}_{\aleph_0}$. (Compare Example 5.1.1.)

Returning to the more noticeable difference between Theorem 39 and Theorem 42, it is actually not obvious that changing 'bounded' to 'closed' is in fact an upgrade. We thus have to put in the work and prove it:

Theorem 43. *Every bounded profile of G is closed.*

Before we endeavour to prove Theorem 43 though, let us first establish Theorem 42, whose proof is straightforward.

Proof of Theorem 42. For each pair Q and Q' of distinguishable profiles in $\mathcal{P}$ let $A_{P,P'} \subseteq \vec{U} = \vec{U}(G)$ be the set of all separations that distinguish P and P' efficiently. Just as in the proof of Theorem 39 the family $\mathcal{A} = (\, A_{P,P'} \mid P \neq P' \,)$ of these sets splinters. To obtain the claimed nested set from Theorem 38 it thus remains to check that each $A_{P,P'}$ is a closed set in $\vec{U}$.

So let P and P' be two distinguishable profiles in $\mathcal{P}$. Let k be the minimum integer for which P and P' induce different profiles on $\vec{S}_k = \vec{S}_k(G)$. Clearly $A_{P,P'} \subseteq \vec{S}_k$, and $\vec{S}_k \subseteq \vec{U}$ is submodular and closed. The restrictions $P \cap \vec{S}_k$ and $P' \cap \vec{S}_k$ are then closed antisymmetric subsets of $\vec{U}$, and since neither P nor P' contains any degenerate element, we have that $A_{P,P'}$ is the set of precisely those $\vec{s} \in \vec{U}$ with $\vec{s} \in P \cap \vec{S}_k$ and $\overleftarrow{s} \in P' \cap \vec{S}_k$ or vice-versa. Therefore, by Proposition 5.2.20, the set $A_{P,P'}$ is closed in $\vec{U}$, concluding the proof. □

Let us now prove Theorem 43. Doing so will demonstrate that Theorem 42 can be applied not only to all (restrictions of) end tangles as seen in Theorem 37, but also to all bounded profiles of G, making it a strictly stronger version of Theorem 39.

For proving Theorem 43 we need one more observation about the structure of bounded profiles. From [18] we know that every closed subset X of $\vec{U}$ has the property that for every infinite chain $\mathcal{C} \subseteq X$ the supremum of $\mathcal{C}$ also lies in X. For a bounded profile to be closed it is therefore necessary to have this property, too. Curiously we shall need this assertion about the existence of suprema, a consequence of being closed, in order to show that bounded profiles are closed in the first place.

Lemma 5.2.21. *Let P be a regular bounded profile in G and $\mathcal{C} = \{(A_i, B_i) \mid i \in I\} \subseteq P$ an infinite increasing chain. Then $(A, B) \in P$ for the supremum (A, B) of $\mathcal{C}$ in $\vec{U}$ given by*

$$A = \bigcup_{i \in I} A_i \qquad \text{and} \qquad B = \bigcap_{i \in I} B_i \,.$$

Proof. Let us start by showing that $(A, B) \in \vec{S}_k$, where k is the integer for which P is a k-profile. For this let $X \coloneqq A \cap B$. Then for each $x \in X$ there is an $(A_i, B_i) \in \mathcal{C}$ with $x \in A_i$. By definition of B we have $x \in B_i$ as well, and hence $x \in A_i \cap B_i$. Since $\mathcal{C}$ is a chain we in fact have $x \in A_j \cap B_j$ for all $j \geqslant i$. We can therefore infer that for every finite subset of X there is some $i \in I$ such that $A_j \cap B_j$ includes that finite subset of X for all $j \geqslant i$. From $|A_j \cap B_j| < k$ for all $j \in I$ it thus follows that $|X| < k$.

We therefore have $(A, B) \in \vec{S}_k$, which means that P orients $\{A, B\}$. If $(A, B) \in P$ we are done, so suppose for a contradiction that $(B, A) \in P$.

By the above observation we may by pass onto a suitable final segment of $\mathcal{C}$ and assume that $X \subseteq (A_i \cap B_i)$ for all $i \in I$. Consider the increasing infinite chain $\mathcal{C}'$ which contains $(A_i \cup B \,,\, B_i \cap A)$ for each $i \in I$. By the assumption that $X \subseteq A_i \cap B_i$ we have that $(A_i \cup B) \cap (B_i \cap A) \subseteq (A_i \cap B_i)$. Hence $\mathcal{C}'$ is a chain in $\vec{S}_k$ as well, and thus $\mathcal{C}' \subseteq P$ by the profile property and $(B, A) \in P$. Let (A', B') be the supremum of $\mathcal{C}'$, that is, let

$$A' = \bigcup_{i \in I} A_i \cup B \qquad \text{and} \qquad B' = \bigcap_{i \in I} B_i \cap A \,.$$

Then clearly $(A', B') = (V, X)$. For each $i \in I$ let C_i be the component of $G - (A_i \cap B_i)$ given by Lemma 5.2.9 for which $(V \setminus C_i \,,\, (A_i \cap B_i) \cup C_i) \in P$. We then have $C_i \subseteq B_i \setminus A_i$. Since $(A \cap B) \subseteq (A_i \cap B_i)$ we in fact have either $C_i \subseteq (B_i \setminus A_i) \cap (B \setminus A)$ or $C_i \subseteq (B_i \setminus A_i) \cap (A \setminus B)$. The former of these would imply $(A, B) \in P$ by consistency, so the latter must hold.

Moreover $C_i \supseteq C_j$ for all $j \geqslant i$, and consequently

$$\bigcap_{i \in I} C_i \subseteq (B' \setminus A') = \emptyset \,.$$

To obtain a contradiction to the boundedness of P using Lemma 5.2.10 it thus remains to verify that $(V \setminus C_i \,,\, C_i \cup N(C_i)) \in P$ for each C_i. But this follows from $N(C_i) \subseteq (A_i \cap B_i)$ and the fact that P is a regular profile, finishing the proof. □

We are now ready to prove that bounded profiles are closed.

Proof of Theorem 43. Let P be any bounded profile in G and k the integer for which P is a k-profile. We shall first deal with the cases where P is irregular, and where we cannot rely on Lemma 5.2.21.

So suppose first that P is irregular. From [15] we then know that $k \leqslant 2$, and that P has a certain form depending on k. If $k = 1$ then $P = \{(V, \emptyset)\}$, which is closed since $\vec{U}$ is Hausdorff. If on the other hand $k = 2$ then P has the form

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

for some $v \in V$ that is not a cutvertex of G. In this case P is closed in $\vec{S}_2$, too: the basic open set of all separations in $\vec{S}_2$ which induce $(\{v\}, \emptyset)$ on $G[\{v\}]$ is a P-avoiding open neighbourhood of all separations in $\vec{S}_2 \setminus P$ except for $(\{v\}, V)$. Since v is not a cutvertex, the latter has as a P-avoiding open neighbourhood the basic open set of all separations which induce $(\{v\}, \{v, w\})$ on $G[\{v, w\}]$, where w is an arbitrary vertex other than v.

Let us now treat the case that P is regular. For this we establish the following:

Claim. *Let (A, B) be a maximal element of P and v any vertex in $B \setminus A$. Further let Y be some subset of $X := A \cap B$. Then no separation in P induces $(X \cup \{v\}, Y)$ on $G[X \cup \{v\}]$.*

To see this, suppose that some $(C, D) \in P$ does induce $(X \cup \{v\}, Y)$ on $G[X \cup \{v\}]$. Consider the supremum $(A \cup C\,,\, B \cap D)$ of (A, B) and (C, D). By assumption the separator of this supremum is a subset of $C \cap D$, and hence $|(A \cup C\,,\, B \cap D)| \leqslant |(C, D)| < k$. By the profile property we thus have $(A \cup C\,,\, B \cap D) \in P$. But this contradicts the maximality of (A, B) in P: since $v \in C \setminus D$ and $v \in B \setminus A$ this supremum is in fact strictly bigger than (A, B). This proves the claim.

We can now show that P is closed. For this let $(A', B') \in P$ be arbitrary; we will construct a basic open neighbourhood of (B', A') in $\vec{S}_k$ which avoids P. From Lemma 5.2.21 and Zorn's Lemma it follows that (A', B') lies below some maximal element (A, B) of P. We set $X := A \cap B$ and fix an arbitrary $v \in B \setminus A$. Let $Y := (A' \cap B') \cap X$. Then (B', A') induces $(X \cup \{v\}, Y)$ on $G[X \cup \{v\}]$. From the claim it follows that the set of all separations which do so is an open neighbourhood of (B', A') which avoids P, hence showing that P is closed. $\square$

Bibliography

[1] A. Berry, R. Pogorelcnik, and G. Simonet, *An introduction to clique minimal separator decomposition*, Algorithms **3** (2010), no. 2, 197–215.

[2] N. Bowler and J. Kneip, *Separations of sets*, Order (2019), 1–15.

[3] J. Carmesin, *A short proof that every finite graph has a tree-decomposition displaying its tangles*, European Journal of Combinatorics **58** (2016), 61–65.

[4] ———, *All graphs have tree-decompositions displaying their topological ends*, Combinatorica **39** (2019), no. 3, 545–596.

[5] 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.

[6] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein, *Connectivity and tree structure in finite graphs*, Combinatorica **34** (2014), no. 1, 11–46.

[7] C. Chekuri and J. Chuzhoy, *Polynomial bounds for the grid-minor theorem*, Journal of the ACM (JACM) **63** (2016), no. 5, 1–65.

[8] D. Coudert and G. Ducoffe, *Clique-decomposition revisited*, INRIA Sophia Antipolis - I3S, 2016.

[9] B. A. Davey and H. A. Priestley, *Introduction to lattices and order*, 2nd ed., Cambridge University Press, 2002.

[10] R. Diestel, *Graph theory (5th edition)*, Springer-Verlag, 2017.
Electronic edition available at `http://diestel-graph-theory.com/`.

[11] R. Diestel, *Ends and tangles*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg, 2017, pp. 223–244.

[12] ———, *Abstract separation systems*, Order **35** (2018), no. 1, 157–170.

[13] ———, *Tree sets*, Order **35** (2018), no. 1, 171–192.

[14] R. Diestel, P. Eberenz, and J. Erde, *Duality theorems for blocks and tangles in graphs*, SIAM Journal on Discrete Mathematics **31** (2017), no. 3, 1514–1528.

[15] ———, *Duality theorems for blocks and tangles in graphs*, SIAM Journal on Discrete Mathematics **31** (2017), no. 3, 1514–1528.

[16] R. Diestel, J. Erde, and D. Weißauer, *Structural submodularity and tangles in abstract separation systems*, Journal of Combinatorial Theory, Series A **167** (2019), 155–180.

[17] R. Diestel, F. Hundertmark, and S. Lemanczyk, *Profiles of separations: in graphs, matroids, and beyond*, Combinatorica **39** (2019), no. 1, 37–75.

[18] R. Diestel and J. Kneip, *Profinite separation systems*, Order (2019), 1–27.

[19] R. Diestel and S.-i. Oum, *Tangle-tree duality in abstract separation systems*, arXiv preprint arXiv:1701.02509 (2017).

[20] ———, *Tangle-tree duality: in graphs, matroids and beyond*, Combinatorica **39** (2019), no. 4, 879–910.

[21] R. Diestel and G. Whittle, *Tangles and the Mona Lisa*, arXiv preprint arXiv:1603.06652 (2016).

[22] C. Elbracht, *Tangles determined by majority vote* (2017).

[23] C. Elbracht, D. Fioravanti, S. Klepper, J. Kneip, L. Rendsburg, M. Teegen, and U. von Luxburg, *Tangles: From weak to strong clustering*, arXiv preprint arXiv:2006.14444 (2020).

[24] C. Elbracht, J. Kneip, and M. Teegen, *Obtaining trees of tangles from tangle-tree duality*, In preparation.

[25] ______, *Tangles are decided by weighted vertex sets*, arXiv preprint arXiv:1811.06821 (2018).

[26] ______, *Trees of tangles in abstract separation systems*, arXiv preprint arXiv:1909.09030 (2019).

[27] ______, *Trees of tangles in infinite separation systems*, arXiv preprint arXiv:2005.12122 (2020).

[28] J. Geelen, *Personal communication*, Oberwolfach.

[29] J P. Gollin and J. Kneip, *Representations of infinite tree-sets*, arXiv preprint arXiv:1908.10327 (2019).

[30] R Halin, *Lattices of cuts in graphs*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg, 1991, pp. 217.

[31] ______, *On the cycle space of an infinite 3-connected graph*, Results in Mathematics **41** (2002), no. 1-2, 106–113.

[32] J. Kneip, *Infinite tree sets and their representations* (2016).

[33] ______, *Ends as tangles*, arXiv preprint arXiv:1909.12628 (2019).

[34] ______, *Profinite tree sets*, arXiv preprint arXiv:1909.12615 (2019).

[35] H.-G. Leimer, *Optimal decomposition by clique separators*, Discrete mathematics **113** (1993), no. 1-3, 99–123.

[36] C.-H. Liu, *Packing and covering immersions in 4-edge-connected graphs*, arXiv preprint arXiv:1505.00867 (2015).

[37] L. Lovász, *Submodular functions and convexity*, Mathematical programming the state of the art, 1983, pp. 235–257.

[38] H. M. MacNeille, *Partially ordered sets*, Trans. Amer. Math. Soc. **42** (1937), no. 3, 416–460. MR1501929

[39] F. Mazoit, *A simple proof of the tree-width duality theorem*, arXiv preprint arXiv:1309.2266 (2013).

[40] N. Robertson and P. D Seymour, *Graph minors. V. Excluding a planar graph*, Journal of Combinatorial Theory, Series B **41** (1986), no. 1, 92–114.

[41] ______, *Graph minors. X. Obstructions to tree-decomposition*, Journal of Combinatorial Theory, Series B **52** (1991), no. 2, 153–190.

[42] N. Robertson, P. D Seymour, and R. Thomas, *Excluding subdivisions of infinite cliques*, Transactions of the American Mathematical Society (1992), 211–223.

[43] ______, *Excluding infinite clique minors*, Vol. 566, American Mathematical Soc., 1995.

[44] R. E Tarjan, *Decomposition by clique separators*, Discrete mathematics **55** (1985), no. 2, 221–232.

[45] A. W. Tucker, *Dual systems of homogeneous linear relations*, 1–18.

[46] K. Wagner, *Über eine Eigenschaft der ebenen Komplexe*, Mathematische Annalen **114** (1937), no. 1, 570–590.

Appendix A

Appendix

A.1 Summary

In this dissertation we study tangles in graphs and in abstract separation systems. Tangles have their roots in graph theory, where they were employed to unify various notions of clusters and connectivity into a single general-purpose framework. Fundamentally tangles differ from previous concepts of connectivity by not imposing certain conditions on some vertex sets, but instead by setting up a system of separations to point towards the cohesive region, following certain consistency axioms.

These separations and their consistency properties can be formulated and studied outside of graph theory as well, making tangles applicable to a wide range of combinatorial structures. Most of the results presented in this work are set in this abstract framework. However, many of our applications will still feature graphs.

This work is divided into three main chapters, Chapters 3 to 5. In Chapter 3 we seek to find out in which ways abstract separation systems differ from known instances of separations such as separations of graphs or bipartitions of sets. For this we provide combinatorial characterisations of those classical types of separations. With these results one can now tell when a given abstract separation system is actually an instance of a well-known type of separations. On the other hand we construct examples of separation systems with a structure that is fundamentally different from any classical type.

We also investigate *submodular* separation systems: those embeddable in some lattice in such a way that they contain either the meet or the join of any two of their elements. This submodularity is used in many places in tangle theory, including in our later chapters. We show that submodularity can be studied without an *a priori* lattice structure, which can then be added *a posteriori* if desired. We also show that certain types of submodular separation systems always admit a separation whose deletion leaves the system submodular.

In Chapter 4 we present results in finite tangle theory. The two archetypes for results in this field are *tree-of-tangles theorems* and *tangle-tree duality theorems*. The first type asserts that if a structure has many tangles, then they can be neatly separated from each other in a tree-like structure. The second says that if a structure has no tangles, then the structure itself has a tree-like shape which

certifies this. We improve upon the state of the art for both types of results.

Regarding tree-of-tangles theorems we identify the core engine of this archetype and prove an elementary and flexible theorem with a compact proof that, as we go on to show, implies virtually all previously known tree-of-tangles theorems. Our results comes with an easy-to-check sufficient condition which makes the derivations of classical results from ours routine. We also use our result to establish tree-of-tangles theorems in settings where this was previously impossible.

As for the tangle-tree duality theorem we improve upon Diestel's and Oum's [19] by giving a new proof of their tangle-tree duality theorem which is not only shorter and conceptually clearer, but also yields a stronger result. In fact we offer two new proofs of this strengthened theorem: both of them follow similar concepts, with the first focusing on brevity and the second on eliminating technical preliminaries. We also show that this tangle-tree duality theorem can be used to establish a tree-of-tangles theorem, demonstrating that these two archetypes of tangle-theoretical results are not as disjoint as they might at first seem. This new technique of constructing a tree of tangles also allows one to obtain bounds on the degrees in that tree, something which previous methods of building trees of tangles are not able to do.

We close Chapter 4 by substantiating the intuition that tangles in graphs point towards clusters or highly cohesive regions. For this we answer a fractional version of a question by Diestel, who asked whether there is for each tangle in a graph a set of vertices such that each separation in the tangle is oriented towards the side containing the majority of those vertices. Using techniques from linear programming we show that this is true with a weighted set of vertices.

Finally, in Chapter 5 we investigate tangles in infinite graphs and abstract separation systems. Every end of an infinite graph defines a tangle of all separations of finite order; we can thus study the properties of ends thorough the tangles they define. The separation system of an infinite graph can be equipped with a natural topology. We give a combinatorial characterisation of those ends which define tangles that are closed in this topology. One of the equivalences we establish calls back to the last part of Chapter 4: the closed end tangles are those decided by majority vote by a finite set of vertices.

We then turn to tree-of-tangle theorems in infinite separation systems. Our core result from Chapter 4 relies on induction and thus fails in an infinite setting. Nevertheless, using the topology from above, we provide a way in which that result can be extended to infinite abstract separation systems. As an application of this we prove a tree-of-tangles theorem for infinite graphs.

A.2 Zusammenfassung

In dieser Dissertation widmen wir uns Tangles in Graphen und in abstrakten Teilungssystemen. Tangles entstammen der Graphentheorie, wo sie eingesetzt wurden um verschiedene Konzepte von Clustern und Zusammenhang zu einem einzigen allgemeinen Ansatz zu vereinen. Tangles unterscheiden sich von bisherigen Zusammenhangskonzepten fundamental dadurch dass sie nicht eine Liste von Bedingungen an gewisse Eckenmengen stellen, sondern stattdessen ein System von Teilungen sind, welche in Richtung der hochzusammenhängenden Region des Graphen deuten und gewissen Konsistenzbedingungen unterliegen.

Diese Teilungen und ihre Konsistenzeigenschaften können auch losgelöst von der Graphentheorie betrachtet und studiert werden, wodurch Tangles auf eine große Bandbreite kombinatorischer Strukturen anwendbar werden. Die meisten der in dieser Arbeit vorgestellten Resultate sind in dieser abstrakten Art formuliert; nichtsdestotrotz sind viele unserer Anwendungen immer noch Graphen.

Die vorliegende Arbeit besteht aus drei Hauptkapiteln, Kapitel 3 bis 5. In Kapitel 3 wollen wir herausfinden inwieweit sich abstrakte Teilungssyteme von den bisher bekannten Arten von Teilungen wie Graphenteilungen oder Bipartitionen unterscheiden können. Hierfür führen wir eine kombinatorische Charakterisierung dieser klassischen Teilunstypen durch. Mit diesen Ergebnissen kann man nun feststellen ob ein gegebenes abstraktes Teilungssystem in Wahrheit eine Instanz eines altbekannten Typs von Teilungen ist. Andererseits konstruieren wir Beispiele von Teilungssystemen mit fundamentalen strukturellen Unterschieden zu den klassischen Art von Teilungen.

Wir untersuchen auch *submodulare* Teilungssyteme: Solche welche so in einen Verband eingebettet werden können, dass sie für je zwei ihrer Elemente jeweils deren Supremum oder Infimum ebenfalls enthalten. Diese Submodularität taucht in der Tangle-Theorie an den verschiedensten Stellen auf und begegnet auch uns in dieser Arbeit mehrmals. Wir zeigen dass Submodularität auch ohne eine *a priori* existierende Verbandsstruktur studiert werden kann, wobei diese Verbandsstruktur nach Wunsch doch noch *a posteriori* hinzugefügt werden kann. Wir zeigen weiterhin dass bestimmte Arten solcher submodularen Teilungssyteme immer eine Teilung enthalten nach deren Löschen das Teilungssystem immer noch submodular ist.

In Kapitel 4 präsentieren wir unsere Ergebnisse in der endlichen Tangle-Theorie. Die zwei Grundtypen von Sätzen hier sind die *Baum-von-Tangles-Sätze* und die *Tangle-Baum-Dualitätssätze*. Erstere besagen dass in einer Struktur mit mehreren Tangles diese in einer baumartigen Weise voneinander getrennt werden können. Letzere Art von Satz besagt dass eine Struktur welche kein Tangle besitzt selbst eine baumartige Form haben muss, an der man dies ablesen kann. Wir präsentieren neue Ansätze für beide Arten von Sätzen.

Für die Baum-von-Tangles-Sätze identifizieren wir die zentrale Wirkungsweise dieser Art von Satz und beweisen damit einen elementaren und flexiblen Satz mit kurzem Beweis welcher, wie wir dann zeigen, nahezu alle bisher existierenden Baum-von-Tangles-Sätze impliziert. Unser Ergebnis nutzt eine einfach zu überprüfende hinreichende Bedingung, wodurch das Ableiten der klassischen Resultate aus unserem zur Routine wird. Wir nutzen unseren Satz auch um neuartige Baum-von-Tangles-Sätze zu beweisen, gerade in Situationen in denen dies bisher nicht möglich war.

Zum Tangle-Baum-Dualitätssatz: Hier verbessern wir Diestels und Oums [19], indem wir einen neuen Beweis für ihren Satz präsentieren, welcher nicht nur kürzer und konzeptuell klarer ist, sondern auch eine stärkere Aussage ermöglicht. Tatsächlich geben wir sogar die Wahl zwischen zwei neuen Beweisen dieses verstärkten Resultats. Beide nutzen ähnliche Konzepte, wobei der erste Beweis auf Kürze ausgelegt ist und der zweite darauf das Prüfen von technischen Vorbedingungen zu eliminieren. Wir zeigen außerdem dass dieser Tangle-Baum-Dualitätssatz genutzt werden kann um einen Baum-von-Tangles-Satz zu beweisen. Die zwei Grundtypen von Sätzen in der Tangle-Theorie sind demnach nicht so verschieden wie es zunächst den Anschein hat. Diese neue Methode um einen Baum von Tangles zu konstruieren ermöglicht es zudem Schranken an die Grade der Ecken in diesem Baum zu erhalten; dies war mit den herkömmlichen Ansätzen zu Bäumen von Tangles nicht möglich.

Wir schließen Kapitel 4 damit ab die Intuition zu untermauern dass Tangles in Graphen auf Cluster oder hochzusammenhängende Regionen zeigen. Hierfür beantworten wir eine fraktionelle Version einer Frage von Diestel: Gibt es für jeden Tangle in einem Graphen eine Menge von Ecken, sodass jede Teilung in dem Tangle auf genau diejenige ihrer Seiten zeigt welche die Mehrheit dieser ausgewählten Ecken enthält? Unter Nutzung von Techniken aus dem linearen Programmieren zeigen wir dass dies mit einer gewichteten Menge von Ecken stimmt.

Schließlich kommen wir in Kapitel 5 zu Tangles in unendlichen Graphen und Teilungssystemen. Jedes Ende eines unendlichen Graphen definiert ein Tangle all seiner Teilungen endlicher Ordnung; wir können also die Eigenschaften von Enden studieren indem wir die von ihnen definierten Tangles analysieren. Das Teilungssystem eines unendlichen Graphen kann auf natürliche Art und Weise mit einer Topologie ausgestattet werden. Wir finden eine kombinatorische Charakterisierung derjenigen Enden deren Tangles in dieser Topologie abgeschlossen sind. Eine der Äquivalenzen, die wir hierbei zeigen, hängt mit dem letzten Teil von Kapitel 4 zusammen: Die abgeschlossenen Tangles von Enden sind genau diejenigen die von einer endlichen Eckenmenge durch Mehrheitsentscheid definiert werden.

Zuletzt widmen wir uns Baum-von-Tangles-Sätzen in unendlichen Teilungssystemen. Unser zentrales Ergebnis aus Kapitel 4 hierzu verwendet Induktion und scheitert daher im Unendlichen. Unter Nutzung der obigen Topolgie können wir dennoch einen Weg aufzeigen wie dieses Ergebnis auf unendliche abstrakte Teilungssysteme erweitert werden kann. Als Anwendung hiervon zeigen wir einen Baum-von-Tangles-Satz für unendliche Graphen.

A.3 Publications related to this thesis

The following (pre-)publications are related to this dissertation:

Chapter 3

Sections 3.1 to 3.3 are based on [2].

Chapter 4

Section 4.1 is based on [26], excepting Sections 4.1.4, 4.1.7, and 4.1.8.
Sections 4.2 and 4.3 are based on [24], excepting Section 4.2.2.
Section 4.4 is based on [25].

Chapter 5

Section 5.1 is based on [33].
Section 5.2 is based on [27], excepting Section 5.2.5.

A.4 Declaration of my contributions

Chapter 3

Sections 3.1 to 3.3 is joint work with Nathan Bowler, who came up with both the original question and first drafts of the proofs of the results therein. The publication [2] was written mostly by myself. In particular the details of Theorem 3's proof, the derivation of Theorem 4, as well as all examples are by me.

Section 3.4 is joint work with Christian Elbracht and Maximilian Teegen. It was their idea to study the class of crossing graphs; in response to this I established Theorem 8, showing that this is simply the class of all graphs. We then proved Theorem 9 together. Theorem 10 is again by me.

Section 3.5 is joint work with Maximilian Teegen, who devised the question and proposed the Dedekind-MacNeille-completion. We then verified together that this completion indeed produces a suitable universe.

Section 3.6 is joint work with Christian Elbracht and Maximilian Teegen, who first discussed the idea of deletable separations in submodular systems with Joshua Erde, and who found the first positive results in Proposition 3.6.1 and Lemma 3.6.2. The theorems 12 and 13 were both proved by me.

Chapter 4

Section 4.1 is joint work with Christian Elbracht and Maximilian Teegen. The very first version of Theorem 18 was discovered by them, back then formulated in terms of profiles and with a significantly more complex proof. While rewriting their proof I realised that involving profiles was necessary in neither the statement nor the proof of Theorem 18. Following this realisation I then also found the much shorter proof that is now presented in Section 4.1.2. The canonical Theorem 19 was developed mainly by Elbracht and Teegen, although our work on canonicity in general was joint. The final version of [26] was written mostly by me.

Regarding the two addendums Sections 4.1.7 and 4.1.8 not included in [26], the first of these is joint and equal work with Christian Elbracht, whereas the latter is my own work.

Section 4.2 is entirely my own work.

Section 4.3 was set off by Nathan Bowler and Joshua Erde, who wondered whether it was possible to obtain a tree-of-tangles theorem by employing the tangle-tree duality theorem in a clever way, but who did not find a way to do so. Christian Elbracht and Maximilian Teegen were the first to find such a method. Following our discussing their workings, on the same day I discovered the proof of Theorem 28 which is presented in this work. (Elbracht's and Teegen's method can be found in [24].) In a presentation of these results Erde then asked whether our approaches could be used to obtain bounds on the degrees in trees-of-tangles; I answered this affirmatively by proving Theorem 29.

Section 4.4 is joint work with Christian Elbracht and Maximilian Teegen. After working on the problem of decider sets for the better part of a year, we finally obtained a working proof of Theorem 33 on the very first day of considering a linear programming approach. This first proof outline was much more involved and included, among other things, an inductive procedure devised by me. With Elbracht's discovery of Lemma 4.4.3 the proof of Theorem 33 became complete and much smoother. Elbracht also constructed the '7-choose-3'-example 4.4.5 which demonstrates that tangles in arbitrary settings may fail to admit a weighted decider; I realised that modifying this to '6-choose-3' also excluded weighted deciders with negative weights. The final version of [25] (excepting its Section 3) was written mostly by me.

Chapter 5

Section 5.1 is entirely my own work.

Section 5.2 is joint work with Christian Elbracht and Maximilian Teegen. The two profinite tree-of-tangles theorems 38 and 40 are both mine. Elbracht and Teegen went on to develop an infinite approach following Section 4.1.4 which yields a canonical tree-of-tangles, something my profinite results are not able to do. Their results can be found in [27]. Of that publication's text, those sections presented here are written predominantly by me. Section 5.2.5 is my own work.

A.5 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.

Jakob Kneip