

Trees of tangles in infinite separation systems

BY CHRISTIAN ELBRACHT, JAKOB KNEIP AND MAXIMILIAN TEEGEN

Fachbereich Mathematik Arbeitsgruppe DM,

Universität Hamburg, Bundesstraße 55, 20146 Hamburg, Germany.

e-mails: christian.elbracht@uni-hamburg.de,

jakob.kneip@uni-hamburg.de, maximilian.teegen@uni-hamburg.de

(Received 14 June 2021; revised 16 June 2021; accepted 14 June 2021)

Abstract

We present infinite analogues of our splinter lemma for constructing nested sets of separations. From these we derive several tree-of-tangles-type theorems for infinite graphs and infinite abstract separation systems.

2020 Mathematics Subject Classification: 05C63 (Primary); 05C83, 05C05,
06A07 (Secondary)

1. Introduction

One of the key concepts in graph minor theory is that of a tree-decomposition of a graph. These allow us to analyse a large graph by dividing it into smaller subgraphs and, working with these, gain insight into the graph as a whole. Tangles, ever since their introduction by Robertson and Seymour [20], have served as a dual object to tree-decompositions of low width in graphs. They have also been studied in their own right, as a way to express, indirectly, highly cohesive substructures such as k -blocks or large grid minors.

Over the years, this theory has been expanded and generalised; the branch of abstract tangle theory emerged. There, the focus shifted from the study of graph minors to tangles as an abstract combinatorial object on an *abstract separation system*, replacing the concrete separations of a graph. These separation systems are just a poset which only expresses how separations lie in relation to each other, keeping little of the information provided by the underlying graph; this notion is abstract enough to be applicable to settings other than graphs. [6–9, 11, 12, 14, 15].

The two central theorems of abstract tangle theory are the *tangle-tree duality theorem* and the *tree-of-tangles theorem*. The former provides a tree-like structure dual to the existence of a tangle, while the latter exposes how tangles can be arranged in a tree structure by way of a nested set of distinguishing separations. When applied to the vertex separations of a graph, such finite nested sets of separations can easily be converted back into a tree-decomposition in the original sense.

In [14] we introduced the ‘splinter lemma’, a unified theorem which implies the known tree-of-tangles theorems for finite separations systems. The merit of this theorem lies in the fact that, while it is strong enough to imply all these results, the proof of the theorem is simple, and its assumptions are easy to check. (See Section 2 for definitions.)

THEOREM 1.1 (Splinter lemma, [14]). *Let U be a universe of separations and $(\mathcal{A}_i)_{i \leq n}$ a family of subsets of U . If $(\mathcal{A}_i)_{i \leq n}$ splinters then we can pick an element a_i from each \mathcal{A}_i so that $\{a_1, \dots, a_n\}$ is nested.*

Theorem 1.1, in a sense, is yet another step in a series of abstractions in the theory of tangles: rather than working with the tangles themselves, it operates just on the collection of sets of separations distinguishing a given pair of these.

Theorem 1.1 is proved by induction: it finds a separation $a_i \in \mathcal{A}_i$ which is nested with some element of every other \mathcal{A}_j , and then proceeds inductively on the remaining $n-1$ family members, restricted to those separations nested with a_i . This approach cannot deal with infinite families of sets, however.

In this paper we overcome these difficulties and present two different ways to obtain a version of Theorem 1.1 for infinite families of sets of separations. Both these versions are as abstract and therefore as widely applicable, as our original Theorem 1.1; they differ in the way they overcome the difficulties of infinite sets of separations, and therefore in the assumptions required.

The first approach is to ensure that an inductive proof adopted from the finite case converges by imposing a closedness condition in a suitable topology: in Section 3 we present a result for profinite separation systems. These constitute a large class of infinite separation systems, including the systems of vertex separations of an infinite graph. In Section 4 we will, as an example, apply this theorem to precisely these graph separations.

However, to deduce the existing tangle-tree theorems in infinite graphs (see [1,3]) we need a second approach which does not rely on these profinite separation systems, as not every tangle of an infinite graph is one of the corresponding profinite separation system. Following this approach, we obtain our second main theorem, which does imply these existing theorems (Theorem 1.4 and 1.5 below) and is therefore a truly infinite analogue of Theorem 1.1. It asks more of the sets \mathcal{A}_i than the profinite version; in return the set of separations we obtain will be canonical, i.e., invariant under isomorphisms:

THEOREM 1.2. *If $(\mathcal{A}_i \mid i \in I)$ thinly splinters with respect to some reflexive symmetric relation \sim on $\mathcal{A} := \bigcup_{i \in I} \mathcal{A}_i$, then there is a set $N \subseteq \mathcal{A}$ which meets every \mathcal{A}_i and is nested, i.e., $n_1 \sim n_2$ for all $n_1, n_2 \in N$. Moreover, this set N can be chosen invariant under isomorphisms: if φ is an isomorphism between (\mathcal{A}, \sim) and (\mathcal{A}', \sim') , then we have $N((\varphi(\mathcal{A}_i) \mid i \in I)) = \varphi(N((\mathcal{A}_i \mid i \in I)))$.*

We prove this statement in Section 5. Like Theorem 1.1, the statement of this theorem is a bit technical, as we want it to be as widely applicable as possible. We then show the usefulness of this abstract theorem throughout Section 6, by deducing the existing theorems about distinguishing tangles in infinite graphs from it.

As a simple example, we start with applying it to tangles in locally finite graphs in Section 6.1. This application is straightforward and demonstrates a prototypical application of Theorem 1.2.

It is also possible to apply Theorem 1.2 to arbitrary infinite graphs, and we do so in Section 6.2. This application uses another new, and interesting, shift of perspective: we cannot apply Theorem 1.2 directly to the sets of separations efficiently distinguishing two profiles since, in general, these do not splinter thinly. Instead we consider a slightly different set, namely, the set of separators, to which Theorem 1.2 does apply:

THEOREM 1.3. *Given a set of distinguishable robust regular profiles \mathcal{P} of a graph G there exists a canonical nested set of separators efficiently distinguishing any pair of profiles in \mathcal{P} .*

This theorem acts as an intermediary result between the existing results about tangles in arbitrary infinite graphs. On the one hand we can, waiving canonicity, transform the nested set of separators back into a nested set of separations, recovering the following result of Carmesin about distinguishing tangles in infinite graphs by way of a nested set of separations:

THEOREM 1.4 ([1, Theorem 5.12]) *For any graph G , there is a nested set N of separations that distinguishes efficiently any two robust principal profiles (that are not restrictions of one another).*

This theorem is a cornerstone in Carmesin's proof that every infinite graph has a tree-decomposition displaying all its topological ends. For more about the relation between ends and tangles also see [4,19]. We deduce Theorem 1.4 from Theorem 1.3 in Section 6.2.1.

On the other hand, if we want to keep canonicity, we can use Theorem 1.3 to deduce a result by Carmesin, Hamann and Miraftab [3]. They construct a canonical object, which they call a tree of tree-decompositions, to distinguish the tangles:

THEOREM 1.5 ([3, Remark 8.3]) *Let G be a connected graph and \mathcal{P} a distinguishable set of principal robust profiles in G . There exists a canonical tree of tree-decompositions with the following properties:*

- (i) *the tree of tree-decompositions distinguishes \mathcal{P} efficiently;*
- (ii) *if $t \in V(T)$ has level k , then (T_t, \mathcal{V}_t) contains only separations of order k ;*
- (iii) *nodes t at all levels have $|V(T_t)|$ neighbours on the next level and the graphs assigned to them are all torsos of (T_t, \mathcal{V}_t) .*

We deduce Theorem 1.5 from Theorem 1.3 in Section 6.2.2.

Theorem 1.3 is also an interesting result in its own right: the set of separators that it provides is a natural intermediate object between the non-canonical nested set of separations in Theorem 1.4 and the canonical tree of tree-decompositions in Theorem 1.5.

Moreover, proving Theorem 1.4 or Theorem 1.5 by first proving Theorem 1.3 and then deducing them breaks up the proof nicely and is, in total, shorter than the original proofs from [1,3].

2. Terminology and basic facts

All terminology for graphs in this paper is from [5]. This paper builds extensively on the framework of [6,14]. A brief summary of their definitions follows, see [6,14] for more in-depth exposition.

A poset \vec{S} together with an order-reversing involution $*$ is called a *separation system*, and its elements are (*oriented*) *separations*. Given an oriented separation \vec{s} we shall denote its *inverse* \vec{s}^* as \overleftarrow{s} and the pair $\{\vec{s}, \overleftarrow{s}\}$ as s . This s is called the *underlying unoriented separation*

of \vec{s} and \overleftarrow{s} , \overleftarrow{s} are its *orientations*. We write S for the set of all underlying unoriented separations of oriented separations in \vec{S} . Informally we refer to either just as ‘separations’ when the distinction is immaterial or the intended meaning is clear.

Two unoriented separation r and s are *nested* if $\vec{r} \leq \vec{s}$ for suitable orientations of r and s . If r and s are not nested we say that they *cross*. Likewise we call oriented separations \vec{r} and \vec{s} *nested* if r and s are nested, and *crossing* otherwise. A *nested set of separations* is one whose elements are pairwise nested.

A separation \vec{s} is *small* if $\vec{s} \leq \overleftarrow{s}$, in which case we say that \overleftarrow{s} is *cosmall*. A separation \vec{s} is *trivial*, if there exists a separation $\vec{r} \in \vec{S}$ such that $r \neq s$ and both, $\vec{s} \leq \vec{r}$ and $\overleftarrow{s} \leq \overleftarrow{r}$. Note that every trivial separation is small. An unoriented separation is said to be *trivial* or *small* if one of its orientations is. A *tree set* in a separation system S is a nested set of separations which does not contain any trivial elements. It is *regular* if it does not contain any small elements.

A separation system \vec{U} is a universe if there are join and meet operators \vee, \wedge which turn the poset into a lattice. A map $f: \vec{U} \rightarrow \vec{U}'$ between universes \vec{U} and \vec{U}' is an *homomorphism of universes* if f commutes with the involutions and the \vee and \wedge operations of \vec{U} and \vec{U}' , i.e. if $f(\overleftarrow{s}) = f(\vec{s})^*$, $f(\vec{r} \vee \vec{s}) = f(\vec{r}) \vee f(\vec{s})$, and $f(\vec{r} \wedge \vec{s}) = f(\vec{r}) \wedge f(\vec{s})$ for all $\vec{r}, \vec{s} \in \vec{U}$. Clearly, if $f: \vec{U} \rightarrow \vec{U}'$ is an homomorphism of universes, then $\vec{r} \leq \vec{s}$ for $\vec{r}, \vec{s} \in \vec{U}$ implies $f(\vec{r}) \leq f(\vec{s})$ in \vec{U}' .

An *isomorphism of universes* is a bijective homomorphism of universes whose inverse is also a homomorphism. Two universes are *isomorphic* if there is an isomorphism between them.

The *corner separations* of two separations s and t are the four separations $\vec{s} \vee \vec{t}$, $\vec{s} \vee \overleftarrow{t}$, $\overleftarrow{s} \vee \vec{t}$, $\overleftarrow{s} \vee \overleftarrow{t}$, or their underlying unoriented separations.

An *order-function* on a universe \vec{U} is some function $|\cdot|: \vec{U} \rightarrow \mathbb{N}_0$ such that $|\vec{s}| = |\overleftarrow{s}| =: |s|$ for every unoriented separation $s \in U$. Note that, unlike in the finite setting, in the infinite setting we require order-functions to take their values in \mathbb{N}_0 rather than in $\mathbb{R}_{\geq 0}$. In the finite case this does not make a big difference, as there were only finitely many orders of separations, which therefore could be scaled to be ‘essentially integer-valued’. Since such an argument is not possible in the infinite set-up, we need to require our order-function to be integer-valued in the first place.

An order-function is *submodular* if for all $\vec{s}, \vec{t} \in \vec{U}$ we have

$$|\vec{s}| + |\vec{t}| \geq |\vec{s} \vee \vec{t}| + |\vec{s} \wedge \vec{t}|.$$

An *orientation* of some set S of unoriented separations is a set $O \subseteq \vec{S}$ containing precisely one of \vec{s} and \overleftarrow{s} for each $s \in S$. An orientation O is said to be *consistent* if there are no $\overleftarrow{r} \leq \vec{s}$ with $r \neq s$ and $\vec{r}, \vec{s} \in O$.

If O_1 and O_2 are two orientations of some (possibly distinct) sets of separations, a separation s *distinguishes* O_1 and O_2 if s has two distinct orientations $\vec{s} \in O_1$ and $\overleftarrow{s} \in O_2$. If an order function is given, and s is of lowest possible order among all separations that distinguish O_1 and O_2 , then s distinguishes them *efficiently*.

For a separation system $\vec{S} \subseteq \vec{U}$ inside some universe \vec{U} , we say that a consistent orientation O of S is a *profile* if it satisfies the profile property:

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

A profile is said to be *robust* if additionally:

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

Note that in [9], Diestel, Hundertmark and Lemanczyk also defined the much more technical but slightly weaker condition of a set of profiles being robust, which only requires the condition of robustness for some separations in U . Every set of distinguishable robust profiles is a *robust set of profiles* in the sense of [9], and on the other hand, the reader familiar with this terminology may replace each appearance of a *set of robust profiles* in this paper by a *robust set of profiles* without changing the proofs.

If U comes with some order function then, for some $k \in \mathbb{N} \cup \{\aleph_0\}$, we write S_k for the set of all separations in U satisfying $|s| < k$. A k -*profile* (in U) is then a profile of S_k and a *profile in U* shall mean a k -profile in U for some k .

Given some universe U and a collection $\mathcal{B} := (\mathcal{A}_i \mid i \in I)$ of subsets \mathcal{A}_i of U , we say that \mathcal{B} *splinters* if it satisfies, for any $i, j \in I$, the following property:

$$\forall s \in \mathcal{A}_i, t \in \mathcal{A}_j: s \in \mathcal{A}_j \text{ or } t \in \mathcal{A}_i \text{ or there exists a corner of } s \text{ and } t \text{ in } \mathcal{A}_i \cup \mathcal{A}_j.$$

One example of a separation system are the (vertex) separations of a graph: given some (infinite) graph G let us consider the separations of G , these are all pairs (A, B) of vertex sets A, B such that $A \cup B = V(G)$ and there is no edge between $A \setminus B$ and $B \setminus A$. For historical reasons and ease of notation the corresponding unoriented separations are represented as $\{A, B\}$ instead of $\{(A, B), (B, A)\}$, contrary to the notation in the abstract set-up. These form a poset where $(A, B) \leq (C, D)$ when $A \subseteq B$ and $C \supseteq D$. The function $*$: $(A, B) \mapsto (B, A)$ is an order-reversing involution. If $A \cap B$ is finite, the order of (A, B) is defined as $|A \cap B|$. Note that the set S_{\aleph_0} of all separations of finite order is a universe of separations where $(A, B) \vee (C, D) = (A \cup C, B \cap D)$ and $(A, B) \wedge (C, D) = (A \cap C, B \cup D)$. A *profile in G* shall be a profile of $S_k \subseteq S_{\aleph_0}$ for some $k \in \mathbb{N} \cup \{\aleph_0\}$.

A profile P in G is *regular* if it does not contain any cosmall separation of G , i.e., it contains no separation of the form $(V(G), X)$. Note that, in graphs, the irregular profiles are not of large interest, since they always point towards either the empty set or a single non-cut-vertex. Formally, we can summarize this statement from [7] as follows:

LEMMA 2.1 ([7]) *Let $G = (V, E)$ be a graph and P an irregular profile in G then either G is connected and $P = \{(V, \emptyset)\}$ or G has a non-cutvertex $x \in V$ such that*

$$P = \{(A, B) \in \vec{S}_2 \mid x \in B \text{ and } (A, B) \neq (\{x\}, V)\}.$$

These irregular profiles are distinguished efficiently from each other and from all other profiles in G by the set of separations

$$\{\{V(G), \emptyset\}\} \cup \{\{V(G), \{x\}\} \mid x \in V(G) \text{ and } x \text{ is not a cutvertex of } G\}.$$

Every separation in this set is nested with all separations of G . Hence, our efforts for applications in graphs will concentrate on regular profiles.

Given some set of vertices $X \subseteq V(G)$, we say that a connected component C of $G - X$ is *tight*, if $N(C) = X$.

For two vertices $x, y \in V(G)$ of a graph G , an x - y -*separator of order k* is a vertex set $X \subseteq V(G) \setminus \{x, y\}$ of size k such that x and y lie in different components of $G - X$. We shall need

the following basic fact about such separators in infinite graphs at various points throughout this paper.

LEMMA 2.2 ([18, Statement 2.4]). *Let G be a graph, $u, v \in V(G)$ and $k \in \mathbb{N}$. Then there are only finitely many separators of size at most k separating u and v minimally.*

Another basic tool is the so-called ‘fish lemma’:

LEMMA 2.3 ([9, Lemma 2.1]) *Let U be a universe and $r, s \in U$ two crossing separations. Every separation t that is nested with both r and s is also nested with all four corner separations of r and s .*

Additionally we shall use the following more general observation about separations nested with a corner separation:

LEMMA 2.4 *Let \vec{r} and \vec{s} be two separations. Every separation nested with one of r or s is also nested with at least one of $\vec{r} \wedge \vec{s}$ and $\vec{r} \vee \vec{s}$.*

Proof. Let t be a separation nested with, say, r . Then t has an orientation \vec{t} with either $\vec{t} \leq \vec{r}$ or $\vec{t} \leq \overleftarrow{r}$. In the first case t is nested with $\vec{r} \vee \vec{s}$ by $\vec{t} \leq \vec{r} \leq (\vec{r} \vee \vec{s})$. In the latter case t is nested with $\vec{r} \wedge \vec{s}$ by $\vec{t} \leq \overleftarrow{r} \leq (\vec{r} \wedge \vec{s})^*$.

The separations that distinguish a given pair of profiles exhibit a lattice-like structure:

LEMMA 2.5 *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.

Moreover we shall need a way to transition between separations and tree-decompositions in graphs. Such a method already exists in finite graphs [2]. The ingredients of that proof together with the results of [17] are all that is needed to show an analogous result for infinite graphs, which we shall present here.

Let us start by recalling the definition of a tree-decomposition. A *tree-decomposition* of a G is a pair (T, \mathcal{V}) of a tree T together with a family $\mathcal{V} = (V_t)_{t \in T}$ of vertex sets $V_t \subseteq V(G)$ such that:

$$(T1) \quad V(G) = \bigcup_{t \in T} V_t;$$

$$(T2) \quad \text{Given } e \in E[G] \text{ there exists a } t \in T \text{ such that } e \subseteq V_t;$$

$$(T3) \quad \text{Given a path } P \text{ in } T \text{ from } t_1 \text{ to } t_3 \text{ and a vertex } t_2 \in P \text{ we have } V_{t_1} \cap V_{t_3} \subseteq V_{t_2}.$$

A separation (A, B) is *induced* by a tree-decomposition (T, \mathcal{V}) if and only if there exists an edge $tt' \in T$ such that for the components $T_t, T_{t'}$ of $T - tt'$ containing t or t' respectively,

we have

$$(A, B) = \left(\bigcup_{t' \in T_t} V_{t'}, \bigcup_{t'' \in T_{t'}} V_{t''} \right).$$

A chain of order type α or a α -chain, or a chain of length α , for some ordinal number α (we shall only use $\alpha \in \mathbb{N}$ or $\alpha = \omega$ or $\alpha = \omega + 1$) is a collection $\{\vec{s}_i \mid 0 \leq i < \alpha\}$ of oriented separations, such that $\vec{s}_i < \vec{s}_j$ whenever $i < j$. A set N of unoriented separations contains such a chain, if all the separations in that chain are orientations of separations from N .

Kneip and Gollin [17] showed the following:

THEOREM 2.6 ([17, Theorem 3.9]) *Every regular tree set which does not contain a chain of order type $\omega + 1$ is isomorphic to the edge tree set of a suitable tree.*

Here the edge tree set of a tree $T = (V, E)$ is the set of separation obtained from the set of oriented edges \vec{E} by ordering them in the natural partial order (that is $v w \in \vec{E}$ is smaller than $x y \in \vec{E}$ if and only if the unique path from v to y meets w and x).

We shall need the following lemma, whose proof is inspired by [2].

LEMMA 2.7 *Let $G = (V, E)$ be an infinite graph and let $N \subseteq S_{\aleph_0}(G)$ be a regular tree set. If we have for every ω -chain $(A_1, B_1) < (A_2, B_2) < \dots$ which is contained in N that $\bigcap_{i \in \mathbb{N}} B_i = \emptyset$, then there exists a tree-decomposition (T, \mathcal{V}) of G whose set of induced separations is \vec{N} .*

Moreover, this tree-decomposition can be chosen canonically: if $\varphi: G \rightarrow G'$ is an isomorphism of graphs, then the tree-decomposition constructed for $\varphi(N)$ in G' is precisely the image under φ of the tree-decomposition constructed for N in G .

Proof. Let $T = (V, E)$ be the tree from Theorem 2.6. Note that by [17, Theorem 3.9(iii)] any isomorphism between the edge tree sets of two distinct trees induces an isomorphism of the underlying trees.

Let α be the isomorphism from the edge tree set of T to N . Given some node $t \in T$ let us denote as F_t the set of oriented separations

$$F_t := \{\alpha(s, t) \mid (s, t) \in \vec{E}\}.$$

We define the bags of our tree-decomposition as $V_t := \bigcap_{(A, B) \in F_t} B$. Let us verify that (T, \mathcal{V}) with $\mathcal{V} = (V_t)_{t \in T}$ is the desired tree-decomposition.

For (T1) let $v \in V$ be given; we need to find a $t \in T$ with $v \in V_t$. If $v \in A \cap B$ for some $(A, B) \in \vec{N}$ then $v \in V_t$ for t being either of the two end-vertices of the edge whose image under α is (A, B) . Otherwise v induces an orientation O of $E(T)$ by orienting each edge $\{x, y\}$ of T as (x, y) if $v \in B \setminus A$ for $(A, B) = \alpha(x, y)$.

Observe that O is consistent. If O has a sink, that is, if there is a node t of T all of whose incident edges are oriented inwards by O , then $v \in V_t$ by definition of O . If O does not have a sink then O contains an ω -chain. This is impossible though, since by definition of O we would have $v \in \bigcap_{i \in \mathbb{N}} B_i$, where (A_i, B_i) is the image under α of the i th element of that ω -chain in O . Thus (T1) holds.

The proof that (T2) holds can be carried out in much the same way due to the fact that every edge of G is included in either A or B for each $(A, B) \in \vec{N}$.

Before we check that (T3) holds, let us show that (T, \mathcal{V}) indeed induces N . For this we need to show that if (x, y) is an oriented edge of T then

$$\alpha(x, y) = \left(\bigcup_{z \in T_x} V_z, \bigcup_{z \in T_y} V_z \right),$$

where T_x and T_y are the components of $T - xy$ containing x and y , respectively. So let $(x, y) \in \vec{E}$ be given and $\alpha(x, y) = (A, B)$. Observe first that $A \cap B \subseteq V_x \cap V_y$ by definition. It thus suffices to show that $A \supseteq \bigcup_{z \in T_x} V_z$ and $B \supseteq \bigcup_{z \in T_y} V_z$ to establish the desired equality.

To see this consider a vertex $v \in V_z$ for some $z \in T_x$. Let \vec{e} be the first edge of the unique $z-x$ -path in T and let $\alpha(\vec{e}) = (A', B')$. We have $\vec{e} \leq (x, y)$ by definition of an edge tree set, and hence $(A', B') \leq (A, B)$ since α is an isomorphism. From this we know that $A' \subseteq A$. We further have $(B', A') \in F_z$ and thus, by definition of V_z , that $v \in A'$. This shows $v \in A$. The argument that $B \supseteq \bigcup_{z \in T_y} V_z$ is similar.

Having established that (T, \mathcal{V}) indeed induces N , we can now deduce from this that (T3) holds: if V_{t_1} and V_{t_3} are two bags of (T, \mathcal{V}) which both contain some vertex v , then v also needs to lie in the separator of every separation that is an image under α of an edge on the path P in T from t_1 to t_3 . Therefore v lies in every V_{t_2} with $t_2 \in P$.

3. The profinite splinter lemma

In this section we establish an extension of the splinter lemma, Theorem 1.1, 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 now.

3.1. Introduction to profinite universes

For an in-depth introduction to profinite separation systems we refer the reader to [10], 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 $p \leq r$ and $q \leq r$ 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 $f_{qp} : X_q \rightarrow X_p$ 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 (cf. [10]).

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 $\vec{r} \leq \vec{s}$ if and only if $\vec{r}_p \leq \vec{s}_p$ 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 an inverse limit of finite universes of separations. The most prominent example of a profinite universe of separations comes from infinite graphs:

EXAMPLE 3.1. Let $G = (V, E)$ be an infinite graph and $\vec{U} = \vec{U}(G)$ the universe of all separations of G , including those (A, B) with $|A \cap B| = \infty$. 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 $f_{ZY} : \vec{U}_Z \rightarrow \vec{U}_Y$ 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 $\mathcal{U} := (\vec{U}_Z \mid Z \in \mathcal{X})$ 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 $f : \vec{U} \rightarrow \varprojlim \mathcal{U}$ 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 $Z := \{v, w\}$ 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 $(\vec{s} \upharpoonright p) := \vec{s}_p$ for the projection of \vec{s} to \vec{U}_p . Likewise, for a set $O \subseteq \vec{U}$ the projection $O \upharpoonright p$ to \vec{U}_p is the set of all $\vec{s} \upharpoonright p$ with $\vec{s} \in O$.

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 3.2 ([10, Lemma 4.1]) *The following statements hold:*

- (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 $\vec{s}_p \in O \setminus p$ 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 [10, Lemma 5.4]:

LEMMA 3.3 *A set $N \subseteq \vec{U}$ is nested if and only if $(N \setminus p) \subseteq \vec{U}_p$ is nested for all $p \in P$.*

A consequence of Lemma 3.2 and Lemma 3.3 (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 3.4 *If $R, S \subseteq \vec{U}$ are closed and $(R \setminus p) \cap (S \setminus p)$ is non-empty for every $p \in P$, then $R \cap S$ is non-empty.*

Proof. The family $((R \setminus p) \cap (S \setminus p) \mid p \in P)$ 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 R and S are closed in \vec{U} , every limit of this family lies in $R \cap S$, which is therefore non-empty.

3.2 Statement and proof of the profinite splinter lemma

Using the framework of profinite separation systems, we can now extend Theorem 1.1 to infinite separation systems:

THEOREM 3.5 *Let $\vec{U} = \varprojlim (\vec{U}_p \mid p \in P)$ be a profinite universe and \mathcal{B} a family of non-empty closed subsets of \vec{U} . If \mathcal{B} splinters then there is a closed 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 $\mathcal{B} \setminus p$ of \mathcal{B} to \vec{U}_p , that is, the family of all $B \setminus p$, 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 $(\vec{r} \setminus p) = \vec{r}_p$ and $(\vec{s} \setminus p) = \vec{s}_p$. Since \mathcal{B} splinters we have either that 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 $(\vec{c} \setminus p) \in B_p \cup B'_p$ 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 1.1 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 non-empty 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 $\vec{\mathcal{U}}$, 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 3.2 and Lemma 3.3. To see that N contains an element of each member of \mathcal{B} , let B be a member of \mathcal{B} . Then $(N \mid p) \cap (B \mid p)$ is non-empty for each $p \in P$ since $(N \mid p) = N_p \in \mathcal{N}_p$, and thus by Lemma 3.4 and the assumption that B is closed in \vec{U} the sets N and B intersect.

4. Application of the profinite splinter lemma

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 3.1, 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 3, we write $(A, B) \mid Z$ 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 3.2, it is easy to observe the following fact about these \vec{S}_k , which will be used throughout this section:

OBSERVATION 4.1 *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 3.5 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 4.2 *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 ([4]) 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 [4]. 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 utilise Theorem 3.5 to establish Theorem 4.2. The latter is a weakening of Theorem 1.4, but our proof via Theorem 3.5 will be significantly shorter and simpler than the proof of Theorem 1.4.

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 4.2 from Theorem 3.5 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 4.3 *Let P and P' be distinguishable regular bounded profiles in G . Then $\mathcal{A}_{P,P'}$ is a closed subset of \bar{U} .*

In order to prove Proposition 4.3 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 [4].

LEMMA 4.4 *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 $(A, B) \leq (V \setminus C_B, C_B \cup N(C_B))$. 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, by definition, $(V \setminus C_B, C_B \cup N(C_B))$ and $(V \setminus C_D, C_D \cup N(C_D))$ in P . 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 $((A, B) \vee (C, D))^* \notin P'$.

Using Lemma 4.4 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 4.5 *Let P and P' be distinguishable regular bounded profiles in G and consider and infinite increasing sequence $(A_1, B_1) \leq (A_2, B_2) \leq \dots$ in $\mathcal{A}_{P,P'}$. Then the sequence of separators $(A_i \cap B_i)_{i \in \mathbb{N}}$ 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 4.4 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, and that $C_i \subseteq B_i$. 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 $(A_i \cap B_i)_{i \geq \mathbb{N}}$. In fact $C_i \supseteq C_j$ for $i \leq j$. 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.

Let us first show that $\bigcap_{i \in \mathbb{N}} C_i$ is empty. If not there is a vertex $v \in \bigcap_{i \in \mathbb{N}} C_i$. By applying Lemma 4.4 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 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. Let us define an orientation \tilde{P} of $S_{\mathbb{N}_0}$ and show that \tilde{P} is a profile extending P , obtaining a contradiction. Let \tilde{P} consist of all $(A, B) \in \tilde{S}_{\mathbb{N}_0}$ 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 defines an orientation of $S_{\mathbb{N}_0}$. Moreover it is easy to see that \tilde{P} is consistent and has the profile property.

To obtain the desired contradiction it is thus left to check that $P \subseteq \tilde{P}$. However any $(A, B) \in P$ with $C_i \subseteq A \setminus B$ for some $i \in \mathbb{N}$ would be inconsistent with $(V \setminus C_i, C_i \cup X_i) \in P$. Therefore \tilde{P} is an extension of P , contrary to the assumption that P is bounded.

Moreover, we can even show that the same statement also holds not only for chains, but even for the entire set $\mathcal{A}_{P,P'}$:

LEMMA 4.6 *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 $(A_1, B_1), (A_2, B_2), \dots$ 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 2.5 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 4.4. By Lemma 4.5 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 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 4.4 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 4.3, i.e. that the sets $\mathcal{A}_{P,P'}$ are closed subsets of \vec{U} :

Proof of Proposition 4.3. Let two distinguishable regular bounded profiles P and P' in G be given. By Lemma 4.6 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 4.4 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 $(A, B) \in \vec{S}_{\mathbb{N}_0}$ 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 3.2. 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 4.2 from Theorem 3.5. Since we shall need a slightly stronger property than splintering at a later point in Section 6, 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 4.7 *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 $|(A, B) \vee (C, D)| \leq |(C, D)|$ or $|(A, B) \vee (D, C)| \leq |(C, D)|$, it follows that this corner separation efficiently distinguishes Q and Q' by Lemma 2.3, 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 4.8 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 2.5 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 2.5 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 4.2:

Proof of Theorem 4.2. By Proposition 4.3, we can apply Theorem 3.5. Thus we only need to show that the collection of these sets $\mathcal{A}_{P,P'}$ splinters. However, this follows from Lemma 4.7 and Lemma 4.8.

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:

EXAMPLE 4.9. Consider the graph displayed in Figure 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 $(2/3) \cdot 2^{n+1}$ which obviously is regular and bounded. The only separation which efficiently distinguishes P_n and P_{n+1} is the separation s_n with separator $\{v_1^i \mid i < n\} \cup \{u_j^{n+1}\}$.




Fig. 1. A locally finite graph where no tree-decomposition distinguishes all the robust regular bounded profiles efficiently. The green separator is the one of the only separation which efficiently distinguishes the profile induced by the K^{64} from the profile induced by the K^{128} .

Additionally, the K^{10} induces a robust profile P_0 of order 4. However the only separation that efficiently distinguishes P_0 and P_1 has the separator $\{v_1^0, w_1^0, w_2^0\}$. But these separations s_1, s_2, \dots , and s_0 can be oriented such as to form a chain of order type $\omega + 1$. This chain witnesses that there cannot be a tree-decomposition which distinguishes all regular bounded profiles efficiently: the separations given by such a tree-decomposition would have to contain this chain of order type $\omega + 1$ which is not possible as every chain in the edge tree set of a tree has length at most ω , cf. Lemma 2.7.

5. The thin splinter lemma

In this section we take a different approach to generalising the finite splinter lemma into an infinite setting. Unlike Theorem 3.5, the result we are going to prove does not require our universe to be ‘closed’. Instead we will require that the separations involved do not, in a sense, cross too badly in that they cross only finitely many separations of lower order.

This will allow us to choose separations that minimise the number of separations crossing them, an idea which also appeared in Carmesin’s original proof of Theorem 1.4 in [1], as well as in [3] and our proof of the canonical splinter lemma for finite separation systems in [14]. However, our theorem here applies to a more general setting and will allow us directly to deduce Carmesin’s theorem for locally finite graphs.

In order to also be able to deduce the full Theorem 1.4 for arbitrary graphs, we will state our theorem in more generality here: not as a theorem about nestedness and separations, but as a theorem about a general nestedness-like relation. This allows us to apply the theorem in Section 6.2 not to separations directly, where it would fail, but to substitute separators as a proxy giving our Theorem 1.3. From this result we will retrieve the separations for our proof of Theorem 1.4 in Section 6.2.1, but we will also build from this a tree of tree-decompositions to deduce Theorem 1.5 in Section 6.2.2.

The statement of our Theorem 1.2 is also inspired by our canonical splinter lemma for the finite setting in [14], and it will too result in a *canonical* nested set, a set which is invariant under isomorphisms.

So let \mathcal{A} be some set and \sim a reflexive and symmetric binary relation on \mathcal{A} . In analogy to our terminology for separation systems, we say that two elements a and b of \mathcal{A} are *nested* if $a \sim b$. Elements of \mathcal{A} that are not nested *cross*. As usual, a subset of \mathcal{A} is nested if all of its elements are pairwise nested, and a single element is nested with a set N if it is nested with every element of N .

In an abuse of notation, given elements a and b of \mathcal{A} , we call $c \in \mathcal{A}$ a *corner* of a and b if every element of \mathcal{A} crossing c also crosses one of a and b . Observe that with this definition corners of elements of \mathcal{A} exhibit the same behaviour as was asserted by Lemma 2.3 for corner separations. However, in contrast to the terminology of separation systems, we do not insist here that a corner of a and b is itself nested with both a and b . This distinction will become relevant in Section 6.2.

Now let $(\mathcal{A}_i \mid i \in I)$ be a family of non-empty subsets of \mathcal{A} and $|\cdot| : I \rightarrow \mathbb{N}_0$ some function, where I is a possibly infinite index set. We shall think of $|i|$ as the order of the elements of \mathcal{A}_i . For an $a \in \mathcal{A}$ and $k \in \mathbb{N}_0$ the *k-crossing number* of a is the number of elements of \mathcal{A} that cross a and lie in some \mathcal{A}_i with $|i| = k$. This *k-crossing number* is either a natural number or infinity. The family $(\mathcal{A}_i \mid i \in I)$ *thinly splinters* if it satisfies the following three properties:

- (i) for every $i \in I$ all elements of \mathcal{A}_i have finite *k-crossing number* for all $k \leq |i|$;
- (ii) if $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$ cross with $|i| < |j|$, then \mathcal{A}_j contains some corner of a_i and a_j that is nested with a_i ;
- (iii) if $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$ cross with $|i| = |j| = k \in \mathbb{N}_0$, then either \mathcal{A}_i contains a corner of a_i and a_j with strictly lower *k-crossing number* than a_i , or else \mathcal{A}_j contains a corner of a_i and a_j with strictly lower *k-crossing number* than a_j .

We are now ready to state and prove the main result of this section:

THEOREM 1.2. *If $(\mathcal{A}_i \mid i \in I)$ thinly splinters with respect to some reflexive symmetric relation \sim on $\mathcal{A} := \bigcup_{i \in I} \mathcal{A}_i$, then there is a set $N \subseteq \mathcal{A}$ which meets every \mathcal{A}_i and is nested, i.e., $n_1 \sim n_2$ for all $n_1, n_2 \in N$. Moreover, this set N can be chosen invariant under isomorphisms: if φ is an isomorphism between (\mathcal{A}, \sim) and (\mathcal{A}', \sim') , then we have $N((\varphi(\mathcal{A}_i) \mid i \in I)) = \varphi(N((\mathcal{A}_i \mid i \in I)))$.*

Proof. We shall construct inductively, for each $k \in \mathbb{N}_0$, a nested set $N_k \subseteq \mathcal{A}$ extending N_{k-1} and meeting every \mathcal{A}_i with $|i| \leq k$, so that the choice of N_k is invariant under isomorphisms. The desired nested set N will then be the union of all these sets N_k .

We set $N_{-1} := \emptyset$. Suppose that for some $k \in \mathbb{N}_0$ we have already constructed a nested set N_{k-1} so that N_{k-1} is canonical and meets every \mathcal{A}_i with $|i| \leq k - 1$. We shall construct a canonical nested set $N_k \supseteq N_{k-1}$ that meets every \mathcal{A}_i with $|i| \leq k$.

Let N_k^+ be the set consisting of the following: for every $i \in I$ with $|i| = k$, among those elements of \mathcal{A}_i that are nested with N_{k-1} , those of minimum *k-crossing number*. We claim that $N_k := N_{k-1} \cup N_k^+$ is as desired.

Since the choice of N_k^+ is invariant under isomorphisms, and N_{k-1} is canonical by assumption, N_k is clearly canonical as well. It thus remains to show that N_k meets every \mathcal{A}_i with $|i| = k$, and that the set N_k is nested.

To see that the former is true, let $i \in I$ with $|i| = k$ be given. It suffices to show that \mathcal{A}_i contains some element that is nested with N_{k-1} . If \mathcal{A}_i already meets N_{k-1} there is nothing to

show, so suppose that it does not. By property (i) every element of \mathcal{A}_i crosses only finitely many elements of N_{k-1} ; pick an $a_i \in \mathcal{A}_i$ that crosses as few as possible. Suppose for a contradiction that a_i crosses some element of N_{k-1} , that is, some $a_j \in \mathcal{A}_j$ with $|j| < |i|$. But then, by property (ii), \mathcal{A}_i contains a corner of a_i and a_j that is nested with a_j . This element of \mathcal{A}_i does not cross a_j and therefore, by virtue of being a corner of a_i and a_j , crosses fewer elements of N_{k-1} than a_i does, contrary to the choice of a_i . Therefore N_k indeed contains an element of each \mathcal{A}_i with $|i| \leq k$.

Let us now show that N_k is nested. Since N_{k-1} is a nested set by assumption, and every element of N_k^+ is nested with N_{k-1} , we only need to show that the set N_k^+ itself is nested. So suppose that some two elements of N_k^+ cross. These two elements then are some $a_i \in \mathcal{A}_i$ and $a_j \in \mathcal{A}_j$ with $|i| = |j| = k$. But now property (iii) asserts that one of \mathcal{A}_i and \mathcal{A}_j contains a corner of a_i and a_j with a strictly lower k -crossing number than the corresponding element a_i or a_j . Since both a_i and a_j are nested with N_{k-1} their corner is nested with N_{k-1} as well, and hence contradicts the choice of a_i or a_j for N_k^+ .

6. Applications of the thin splinter lemma

In this section we are going to apply Theorem 1.2 to infinite graphs. The application to locally finite graphs in Section 6.1 will be a straightforward application to a universe of separations, whereas in Section 6.2 we are going to use a more involved argument.

For either case we will utilise the fact that separations which efficiently distinguish two regular profiles are tight. Recall that for a set $X \subseteq V$ a component C of $G - X$ is *tight* if $N(C) = X$. We say that a separation (A, B) of G is *tight* if for $X := A \cap B$ each of $A \setminus B$ and $B \setminus A$ contains some tight component of $G - X$.

LEMMA 6.1 *Let P, P' be two distinct regular profiles in an arbitrary graph G . If (A, B) is a separation of finite order that efficiently distinguishes P and P' , then (A, B) is tight.*

Proof. Let $(A, B) \in P, (B, A) \in P'$.

Suppose for a contradiction that $B \setminus A$ does not contain a tight component of $G - (A \cap B)$. Let Y_1, \dots, Y_m , be an enumeration of all proper subsets of $A \cap B$. For every Y_l let \mathcal{C}_l be the set of components of $G - (A \cap B)$ in B with neighbourhood exactly Y_l . By consistency of P' we have $(\bigcup \mathcal{C}_l \cup Y_l, V \setminus \bigcup \mathcal{C}_l) \in P'$. Since moreover (A, B) efficiently distinguishes P from P' and $|Y_l| < |A \cap B|$ we know that $(\bigcup \mathcal{C}_l \cup Y_l, V \setminus \bigcup \mathcal{C}_l) \in P$ as well. Moreover, $(A \cap B, V) \in P$ since P is regular. Thus, by an inductive application of the profile property P we have that for every l

$$(A \cap B, V) \vee \left(\bigcup \mathcal{C}_1 \cup Y_1, V \setminus \bigcup \mathcal{C}_1 \right) \vee \dots \vee \left(\bigcup \mathcal{C}_l \cup Y_l, V \setminus \bigcup \mathcal{C}_l \right) \in P.$$

However, for $l = m$ this contradicts the assumption since

$$(A \cap B, V) \vee \left(\bigcup \mathcal{C}_1 \cup Y_1, V \setminus \bigcup \mathcal{C}_1 \right) \vee \dots \vee \left(\bigcup \mathcal{C}_m \cup Y_m, V \setminus \bigcup \mathcal{C}_m \right) = (B, A) \notin P.$$

6.1. Locally finite graphs

In this section we apply Theorem 1.2 to the set of separations of a locally finite graph, which will result in a canonical nested set of separations efficiently distinguishing any two distinguishable regular profiles in G . The proof of this theorem will be a straightforward application of Theorem 1.2 to sets $\mathcal{A}_{P, P'}$ of separations efficiently distinguishing two profiles

in G . Following the strategy of this proof, one might be able to obtain similar results for other infinite separation systems, e.g., in a matroid.

So let $G = (V, E)$ be a locally finite connected graph and \mathcal{P} a set of robust regular profiles in G .

Let I be the set of pairs of distinguishable profiles in \mathcal{P} . For each pair P and P' of distinguishable profiles in \mathcal{P} let $\mathcal{A}_{P,P'}$ be the set of all separations of G that distinguish P and P' efficiently. Observe that by definition all separations in $\mathcal{A}_{P,P'}$ are of the same order; let us write $|P, P'|$ for this order.

Let \mathcal{A} be the union of all the $\mathcal{A}_{P,P'}$. We wish to show that $(\mathcal{A}_i \mid i \in I)$ thinly splinters, using as the relation \sim on \mathcal{A} the usual nestedness of separations. We shall prove first that property (i) is satisfied, i.e. that each separation in an $\mathcal{A}_{P,P'}$ crosses only finitely many other separations from sets $\mathcal{A}_{Q,Q'}$ with $|Q, Q'| \leq |P, P'|$.

Making use of the tightness of the separations in the $\mathcal{A}_{P,P'}$, property (i) will follow immediately from the following assertion:

PROPOSITION 6.2 *Let (A, B) be a separation that efficiently distinguishes some two regular profiles in G . Then G has only finitely many tight separations of order at most $|(A, B)|$ that cross (A, B) .*

We shall derive Proposition 6.2 from the following lemma about tight separations:

LEMMA 6.3 *Let (A, B) and (A', B') be two tight separations of G . Then (A', B') is either nested with (A, B) , or its separator $A' \cap B'$ is a \subseteq -minimal x - y -separator in G for some pair x, y of vertices from $(A \cap B) \cup N(A \cap B)$.*

Proof. Since (A', B') is tight each of $A' \setminus B'$ and $B' \setminus A'$ contains some tight component of $G - (A' \cap B')$. If $A \cap B$ meets all tight components of $G - (A' \cap B')$ then in particular $A \cap B$ meets these two components, say in x and in y . But then $A' \cap B'$ is a \subseteq -minimal x - y -separator with $x, y \in A \cap B$.

Therefore we may assume that $A \cap B$ misses some tight component C' of $G - (A' \cap B')$. By switching their names if necessary we may assume that this component C' is contained in $A \setminus B$. Since $C' \subseteq A \setminus B$ has no neighbours in $B \setminus A$ but has $A' \cap B'$ as its neighbourhood we can infer that $(A' \cap B') \subseteq A$.

Consider now a tight component C of $G - (A \cap B)$ that is contained in $B \setminus A$. From $(A' \cap B') \subseteq A$ it follows that C does not meet $A' \cap B'$ and is hence contained in either $A' \setminus B'$ or $B' \setminus A'$. By possibly switching the roles of A' and B' we may assume that $C \subseteq A' \setminus B'$. As above we can conclude from the tightness of C that $(A \cap B) \subseteq C$.

It remains to check two cases. If $(B \setminus A) \cap (B' \setminus A')$ is empty we have $B \subseteq A'$ and $B' \subseteq A$, that is, that (A', B') is nested with (A, B) . The other remaining case is that $(B \setminus A) \cap (B' \setminus A')$ is non-empty.

In that case, since G is connected, the set $(A \cap B) \cap (A' \cap B')$ must be non-empty as well, since $N((B \setminus A) \cap (B' \setminus A')) \subseteq (A \cap B) \cap (A' \cap B')$. Pick a vertex z from that set. Since (A', B') is tight z has neighbours x and y in some tight components of $G - (A' \cap B')$ contained in $A' \setminus B'$ and in $B' \setminus A'$, respectively. Then $A' \cap B'$ is a \subseteq -minimal x - y -separator in G , and moreover $x, y \in (A \cap B) \cup N(A \cap B)$ since $z \in A \cap B$.

Let us now use Lemma 6.3 to establish Proposition 6.2:

Proof of Proposition 6.2. Since G is locally finite the set $(A \cap B) \cup N(A \cap B)$ is finite. Therefore, by Lemma 2.2, there are only finitely many \subseteq -minimal $x - y$ -separators of size at most $|A, B|$ with $x, y \in (A \cap B) \cup N(A \cap B)$. Leveraging again the fact that G is locally finite and using that G is connected, we get that there are only finitely many separations of G with such a separator.

The assertion now follows from Lemma 6.3 since we know by Lemma 6.1 that (A, B) is tight.

The family $(\mathcal{A}_i \mid i \in I)$ therefore has property (i). With regard to property (ii) it turns out that we already did the required work back in Section 4:

COROLLARY 6.4 *If $\vec{a}_i \in \mathcal{A}_i$ and $\vec{a}_j \in \mathcal{A}_j$ cross with $|i| < |j|$, then \mathcal{A}_j contains some corner separation of a_i and a_j .*

Proof. This is the assertion of Lemma 4.7.

It remains to show that $(\mathcal{A}_i \mid i \in I)$ has property (iii).

Using again our preparatory work from Section 4, we can now show that $(\mathcal{A}_i \mid i \in I)$ has property (iii) using Lemma 2.4:

LEMMA 6.5 *If $\vec{a}_i \in \mathcal{A}_i$ and $\vec{a}_j \in \mathcal{A}_j$ cross with $k = |i| = |j|$, then either \mathcal{A}_i contains a corner separation of a_i and a_j with strictly lower k -crossing number than \vec{a}_i , or else \mathcal{A}_j contains a corner separation of a_i and a_j with strictly lower k -crossing number than \vec{a}_j .*

Proof. By switching their roles if necessary we may assume that the k -crossing number of \vec{a}_i is at most the k -crossing number of \vec{a}_j .

From Lemma 4.8 it follows that \mathcal{A}_j contains a corner separation of a_i and a_j whose opposite corner separation lies in either \mathcal{A}_i or \mathcal{A}_j . Now Lemma 2.4 implies that the sum of the k -crossing numbers of this pair of opposite corner separations is at most the sum of the k -crossing numbers of \vec{a}_i and \vec{a}_j . This inequality is in fact strict since \vec{a}_i and \vec{a}_j cross each other but are each nested with both corner separations.

If the first corner separation is not already as desired, that is, if its k -crossing number is not strictly lower than the k -crossing number of \vec{a}_j , we can infer that the k -crossing number of the opposite corner separation is strictly lower than that of \vec{a}_i . Since we assumed in the beginning that the k -crossing number of \vec{a}_i is no greater than that of \vec{a}_j this proves the claim.

We are now ready to prove the main result of this subsection, which is similar to [3, Theorem 7.5]:

THEOREM 6.6 *Let G be a locally finite connected graph and \mathcal{P} some set of robust regular profiles in G . Then there exists a nested set \mathcal{N} of separations which efficiently distinguishes any two distinguishable profiles in \mathcal{P} . Moreover, this set is canonical, i.e. invariant under isomorphisms: If $\alpha: G \rightarrow G'$ is an isomorphism, then $\alpha(\mathcal{N}(G, \mathcal{P})) = \mathcal{N}(\alpha(G), \alpha(\mathcal{P}))$.*

Proof. The combination of Proposition 6.2, Corollary 6.4 and Lemma 6.5 shows that the family $(\mathcal{A}_i \mid i \in I)$ thinly splinters. The nested set $N \subseteq \mathcal{A}$ produced by Theorem 1.2 meets each set \mathcal{A}_i and thus distinguishes all pairs of distinguishable profiles in \mathcal{P} efficiently.

The nested set found by Theorem 6.6 does not in general correspond to a tree-decomposition of G , as Example 4.9 demonstrated. However Theorem 6.6 can be used to show that for every fixed integer k the subset of \mathcal{N} consisting of all separations of order at

most k gives rise to a tree-decomposition of G , as this subset will satisfy the conditions from Lemma 2.7. In particular we can use Theorem 6.6 together with Lemma 2.7 to prove [3, Theorem 7.3], that there is for every $k \in \mathbb{N}$, every locally finite graph G and every set \mathcal{P} of distinguishable robust regular profiles, pairwise distinguishable by a separation of order at most k , a canonical tree-decomposition of G that efficiently distinguishes all profiles from \mathcal{P} .

6.2. Graphs with vertices of infinite degree

When we consider graphs with vertices of infinite degree, the method of the previous section fails as we lose Proposition 6.2: It does not necessarily hold that every separation in an $\mathcal{A}_{P,P'}$ crosses only finitely many other separations from sets $\mathcal{A}_{Q,Q'}$ with $|Q, Q'| \leq |P, P'|$. Moreover, Dunwoody and Krön [13] gave an example of a graph which does not contain a canonical nested set of separations separating its ends. As ends induce robust regular profiles, in arbitrary graphs, it is not generally possible to find a canonical nested set of separations distinguishing all the robust regular profiles.

To show the result for locally finite graphs we made use of the observation that only finitely many different *separators* are involved, and then used that every separator appears in only finitely many separations. Thus in this section instead of applying Theorem 1.2 directly to some set of *separations*, we are going to apply it to only the set of *separators*.

With this approach we show that in an arbitrary graph you can find a canonical nested set of separators which efficiently distinguishes all the robust regular profiles in G . We shall make the meaning of this more precise shortly. We propose that this set of separators is a natural intermediate object for distinguishing profiles. Moreover we will show that if we restrict ourselves to the set of robust principal profiles – which we will define at the end of this section – then from this set we can build both a non-canonical nested set of separations as in Theorem 1.4 (from [1]) as well as a canonical tree of tree-decompositions in the sense of [3].

Either of these objects can trivially be converted back to a set of separators. Our technique splits the process of building either of these cleanly into two independent steps, which makes it more accessible than the proofs in [1] and [3]. Moreover, the first step of this process also works for non-principal but regular profiles, allowing us to also get an (intermediate) result for those profiles, unlike the theorems from [1] and [3]. Note that distinguishing non-principal profiles is also discussed extensively in [16].

Many of the techniques applied throughout are similar to or inspired by arguments made in [3], particularly the approach of minimising the crossing-number, even though the different levels of abstraction make it hard to draw concrete parallels.

Let us now begin with the formal notation. We say that a *set of vertices* $X \subseteq V(G)$ *efficiently distinguishes* a pair P and P' of profiles in G if there exists a separation (A, B) of G with separator $A \cap B = X$ which efficiently distinguishes P and P' . Such a separation (A, B) is then a *witness* that X efficiently distinguishes P and P' .

Given some set of distinguishable robust regular profiles \mathcal{P} of an (infinite) graph G , we define as \mathcal{A} the set of all such separators X which distinguish some pair of profiles in \mathcal{P} efficiently. We say that a separator X is *nested* with $Y \in \mathcal{A}$, i.e. $X \sim Y$, whenever X is contained in $C \cup Y$ for some component C of $G - Y$. In other words Y does not properly separate any two vertices of X . This relation is reflexive, the following lemma shows that it is also symmetric on \mathcal{A} . Unfortunately, its natural extension to all finite subsets of $V(G)$ is not. The reader should take note that this will lead to some situations where we argue that some set Y is nested with some $X \in \mathcal{A}$ *provided that* $Y \in \mathcal{A}$.

LEMMA 6.7 *If $X, Y \in \mathcal{A}$ and X is contained in Y together with some component of $G - Y$, then Y is contained in X together with some component of $G - X$.*

Proof. Pick a separation (A, B) witnessing that $X \in \mathcal{A}$. Since this separation efficiently distinguishes two regular profiles, by Lemma 6.1, there are at least two tight components of $G - X$, one in either side of (A, B) . At least one of these tight components, say C , does not meet Y and is therefore contained in a connected component C' of $G - Y$. Now, as required, we find

$$X = N(C) \subseteq C \cup N(C) \subseteq C' \cup N(C') \subseteq C' \cup Y.$$

As usual we take as I the set of pairs of distinguishable profiles in \mathcal{P} . But this time we define $\mathcal{A}_{P,P'}$ for each pair P, P' in I to be the set of all the sets of vertices in G which distinguish P and P' efficiently. All these separators in $\mathcal{A}_{P,P'}$ have the same size; this size shall be $|P, P'|$.

We claim that $\{\mathcal{A}_{P,P'} \mid \{P, P'\} \in I\}$ thinly splinters. Before we can show property (i) we need to make two basic observations about how the vertices of a crossing pair of separators in \mathcal{A} lie:

LEMMA 6.8 *If $X, Y \in \mathcal{A}$ cross, then Y contains a vertex from every tight component of $G - X$.*

Proof. If C is a tight component of $G - X$ such that Y does not contain any vertex of C then C is contained in some component C' of $G - Y$. However, then $X = N(C) \subseteq C' \cup Y$, i.e., X is nested with Y contradicting the assertion.

LEMMA 6.9 *If $X, Y \in \mathcal{A}$ cross, then Y contains a pair of vertices v and w such that X is a \subseteq -minimal v - w -separator.*

Proof. There are at least two tight components C_1, C_2 of $G - X$ and Y meets both of them by Lemma 6.8. Let v be a vertex in $Y \cap C_1$ and w a vertex in $Y \cap C_2$. As both C_1 and C_2 are tight components, X is indeed a \subseteq -minimal v - w -separator.

We can now combine these with Lemma 2.2 to prove property (i).

LEMMA 6.10 *For every pair of profiles $P, P' \in \mathcal{P}$ every $X \in \mathcal{A}_{P,P'}$ has finite k -crossing-number for all $k \leq |P, P'|$.*

Proof. By Lemma 6.9, for every $Y \in \mathcal{A}$ of size k which crosses X , there are vertices $v, w \in X$ which are minimally separated by Y . However, there is only a finite number of pairs of vertices v, w in X and by Lemma 2.2 every pair has only finitely many minimal separators of size k . Therefore only finitely many such $Y \in \mathcal{A}$ exist.

The following lemmas show how the separators of corner separations behave under our new nestedness relation. We will need these to prove properties (i) and (ii). Recall from Section 5 that a *corner* of two separators $X, Y \in \mathcal{A}$ is a separator $Z \in \mathcal{A}$ which crosses only elements of \mathcal{A} which cross either X or Y . Note that this does not imply that Z is nested with X and Y .

LEMMA 6.11 *Let $X, Y \in \mathcal{A}$ be a crossing pair of separators and let (A_X, B_X) and (A_Y, B_Y) , respectively, be separations which witness that these are in \mathcal{A} . Then for every*

$Z \in \mathcal{A}$ which is nested with both X and Y there is a component C_Z of $G - Z$, such that $X \cup Y \subseteq C_Z \cup Z$. In particular $(A_X \cup A_Y) \cap (B_X \cap B_Y)$, the separator of $(A_X, B_X) \vee (A_Y, B_Y)$, is a corner of X and Y provided that it lies in \mathcal{A} .

Proof. We first show that Z does not separate X and Y . Since Z is nested with X and X efficiently distinguishes two regular profiles there is, by Lemma 6.1, a tight component C_X of $G - X$ which is disjoint from Z . By Lemma 6.8, there is a vertex $y \in C_X \cap Y \subseteq Y \setminus Z$.

By a symmetrical argument there also exists a vertex $x \in X \setminus Z$. Since C_X is tight there is a path from x to y contained in C_X except for x . This path avoids Z .

Now, since Z is nested with X there is a component C_Z of $G - Z$ which contains $X \setminus Z$. In particular this component contains x . Similarly, there is a component of $G - Z$ containing $Y \setminus Z$ and hence, in particular, y . Since Z does not separate x and y this component is the same as C_Z . Therefore $X \cup Y \subseteq C_Z \cup Z$, as required. In particular, if $(A_X \cup A_Y) \cap (B_X \cap B_Y) \in \mathcal{A}$ then $(A_X \cup A_Y) \cap (B_X \cap B_Y) \subseteq C_Z \cup Z$, hence $(A_X \cup A_Y) \cap (B_X \cap B_Y) \sim Z$ and therefore $(A_X \cup A_Y) \cap (B_X \cap B_Y)$ is a corner of X and Y .

LEMMA 6.12 *Let $X, Y \in \mathcal{A}$ be a crossing pair of separators and let (A_X, B_X) and (A_Y, B_Y) , respectively, be witnesses that these are in \mathcal{A} . If $Z \in \mathcal{A}$ is nested with X , and each of the corner separations $(A_X, B_X) \vee (A_Y, B_Y)$ and $(A_X, B_X) \wedge (A_Y, B_Y)$ distinguishes some pair of profiles efficiently then Z is nested with one of the separators $(A_X \cup A_Y) \cap (B_X \cap B_Y)$ or $(A_X \cap A_Y) \cap (B_X \cup B_Y)$.*

Proof. Since Z and X are nested there is a component C^Z of $G - X$ such that $Z \subseteq C^Z \cup X$. Let us assume without loss of generality that $C^Z \subseteq A_X$, we will show that Z is nested with $(A_X \cup A_Y) \cap (B_X \cap B_Y)$.

Since $(A_X, B_X) \vee (A_Y, B_Y)$ efficiently distinguishes some regular profiles there is, by Lemma 6.1, a tight component of $(A_X \cup A_Y) \cap (B_X \cap B_Y)$ contained in $(B_X \cap B_Y)$. However $Z \subseteq A_X$, so this component cannot meet Z . Hence, by Lemma 6.8, Z cannot cross the separator $(A_X \cup A_Y) \cap (B_X \cap B_Y)$.

These now allow us to reuse Lemmas 4.7 and 4.8 to prove properties (ii) and (iii):

LEMMA 6.13 *If two separators $X \in \mathcal{A}_{P,P'}$ and $Y \in \mathcal{A}_{Q,Q'}$ cross and $|P, P'| < |Q, Q'|$ then there is a corner $Y' \in \mathcal{A}_{Q,Q'}$ of X and Y which is nested with X .*

Proof. Let (A_X, B_X) be a separation witnessing that $X \in \mathcal{A}_{P,P'}$ and let (A_Y, B_Y) be a separation witnessing that $Y \in \mathcal{A}_{Q,Q'}$. By Lemma 4.7 there is a corner separation of (A_X, B_X) and (A_Y, B_Y) which also distinguishes Q and Q' efficiently. The separator Y' of this corner separation does not meet all tight components of $G - X$, so Y' is nested with X and thus is by Lemma 6.11 as desired.

LEMMA 6.14 *If two separators $X \in \mathcal{A}_{P,P'}$ and $Y \in \mathcal{A}_{Q,Q'}$ cross and $|P, P'| = |Q, Q'| = k$ then either there is a corner $Y' \in \mathcal{A}_{Q,Q'}$ of X and Y which has a strictly lower k -crossing-number than Y , or there is a corner $X' \in \mathcal{A}_{P,P'}$ of X and Y which has strictly lower k -crossing-number than X .*

Proof. By switching their roles if necessary we may assume that the k -crossing number of Y is at most the k -crossing number of X . Let (A_X, B_X) be a separation witnessing that $X \in \mathcal{A}_{P,P'}$ and let (A_Y, B_Y) be a separation witnessing that $Y \in \mathcal{A}_{Q,Q'}$. By Lemma 4.8 there

is a corner separation of (A_X, B_X) and (A_Y, B_Y) which efficiently distinguishes P and P' and whose opposite corner separation efficiently distinguishes either P and P' or Q and Q' . Let us denote their separators as Z and Z' respectively.

By Lemmas 6.11 and 6.12 and the fact that Z and Z' are nested with both X and Y we have that Z and Z' are corners of X and Y and that the sum of the k -crossing numbers of Z and Z' is strictly lower than the sum of the k -crossing numbers of X and Y .

Thus, if the k -crossing number of Z is strictly lower than the k -crossing number of X , we can take Z for X' . Otherwise we can infer that the k -crossing number of Z' is strictly lower than that of Y . Since we assumed in the beginning that the $|i|$ -crossing number of Y is not greater than that of X . This proves the claim since we can then take Z' for X' or Y' , depending.

With this all the requirements of Theorem 1.2 are satisfied. Immediately we obtain the main result of this section:

THEOREM 1.3. *Given a set of distinguishable robust regular profiles \mathcal{P} of a graph G there exists a canonical nested set of separators efficiently distinguishing any pair of profiles in \mathcal{P} .*

As noted before, to be able to deduce Theorem 1.4 and [3, Remark 8.3] we restrict our set \mathcal{P} to be a set of *principal* robust profiles. A k -profile P in G is *principal* if it contains for every set X of less than k vertices a separation of the form $(V(G) \setminus C, C \cup X)$ where C is a connected component of $G - X$. In particular, every principal profile is regular. Note that this notion of principal profiles is equivalent to the notion of ‘profiles’ in Carmesin’s [1]; the term *principal profiles* comes from [3]. Observe that in locally finite graphs an inductive application of the profile property P shows that every profile is principal.

This restriction to principal profiles is necessary for Theorem 1.4, as Elm and Kurkofka [16, Corollary 3.4] have shown that there is a graph together with a set of (non-principal but robust and distinguishable) profiles, which do not permit the existence of a nested set of separations distinguishing all of them.

6.2.1. Nested sets of separations

If we restrict \mathcal{P} to a set of principal profiles, the nested set of separators from Theorem 1.3 can be transformed into a nested set of *separations* which still distinguishes all the profiles in \mathcal{P} if we give up on canonicity. This task is not entirely trivial.

The natural approach would be to take for each separator every one of the separations belonging to one of its tight components, i.e. the separation $(C \cup X, V \setminus C)$ for every tight component C of $G - X$. However, if the separators overlap the resulting set of separations might not be nested. The following recent result by Elm and Kurkofka states that we need to omit no more than one of the tight components for each separator to reclaim nestedness.

THEOREM 6.15 ([16, Corollary 6.1]) *Suppose that \mathcal{Y} is a principal collection of vertex sets in a connected graph G . Then there is a function \mathcal{K} assigning to each $X \in \mathcal{Y}$ a subset $\mathcal{K}(X) \subseteq \mathcal{C}_X$ (the set \mathcal{C}_X consists of the components of $G - X$ whose neighbourhoods are precisely equal to X) that misses at most one component from \mathcal{C}_X , such that the collection*

$$\{\{V \setminus K, X \cup K\} \mid X \in \mathcal{Y} \text{ and } K \in \mathcal{K}(X)\}$$

is nested.

Here, a principal collection of vertex sets is just a set \mathcal{Y} of subsets of V such that, for every $X, Y \in \mathcal{Y}$, there is at most one component of $G - X$ which is met by Y . In particular, any nested set of separators is a principal collection of vertex sets.

Having for every separator all but one of these tight component separations is still enough to efficiently distinguish all the profiles in \mathcal{P} . However, as Theorem 6.15 does not give as a canonical choice for the function \mathcal{K} , we need to give up the canonicity at this point. However, this still allows us to prove the following theorem by Carmesin:

THEOREM 1.4 ([1, Theorem 5.12]) *For any graph G , there is a nested set N of separations that distinguishes efficiently any two robust principal profiles (that are not restrictions of one another).*

Proof. If G is not connected, then every robust principal profile of G induces a robust principal profile on exactly one of the connected components of G . It is easy to see that we can then apply the theorem to all connected components from G independently and obtain our desired nested set of separations of G from those of the connected components together with separations of the form $(C, V \setminus C)$ for connected components C of G . Thus let us suppose that G is connected.

Let N be the nested set of separations obtained by applying Theorem 6.15 to the set \mathcal{N} of separators obtained from Theorem 1.3. Given any two profiles $P, Q \in \mathcal{P}$ there is a separator X in \mathcal{N} which efficiently distinguishes P and Q . By Lemma 6.1 there are two distinct tight components C and C' of $G - X$ such that both $(V \setminus C, C \cup X) \in P$ and $(C' \cup X, V \setminus C') \in P$ efficiently distinguish P and Q . However, at least one of these two separations is an element of N .

For the reader's convenience, we also offer a direct proof of Theorem 1.4 which does not use Theorem 6.15. Instead we perform an argument akin to one of the arguments used in the proof of Theorem 6.15 but in slightly simpler form, as the statement we need is a weaker one than Theorem 6.15.

Direct proof of Theorem 1.4. Let \mathcal{N} be the nested set of separators obtained from Theorem 1.3 applied to the set of robust principal profiles. Pick an enumeration of \mathcal{N} which is increasing in the size of the separators, i.e., an enumeration $\mathcal{N} = \{X_\alpha \mid \alpha < \beta\}$ such that $|X_\alpha| \leq |X_\beta|$ whenever $\alpha < \beta$.

We will construct a transfinite ascending sequence of nested sets $(N_\gamma)_{\gamma \leq \beta}$, of separations. Each N_γ will contain only separations with separators in $\{X_\alpha \mid \alpha < \gamma\}$, and every pair of profiles efficiently distinguished by such a separator X_α , $\alpha < \gamma$, will also be efficiently distinguished by some separation in N_γ .

For the successor steps of our construction suppose that we already constructed N_γ and consider X_γ . Since X_γ is nested with all X_α satisfying $\alpha < \gamma$ we know that X_γ induces a consistent orientation of N_γ since any separation $(A, B) \in N_\gamma$ satisfies either $X_\gamma \subseteq A$ or $X_\gamma \subseteq B$ but not both, as $|(A, B)| \leq |X_\gamma|$.

Consider the set \mathcal{C} of tight components of $G - X_\gamma$ and let \mathcal{D} be the set of the remaining, non-tight, components of $G - X_\gamma$.

Given any separation $(A, B) \in N_\gamma$ pointing away from X_γ (that is $X_\gamma \subseteq A$), the side B is contained in the union of one component $C_B \in \mathcal{C}$ together with some components in \mathcal{D} : Since X_γ is nested with $A \cap B$ there is a component in $G - X_\gamma$ containing $(A \cap B) \setminus X_\gamma$, thus, any other component C of $G - X_\gamma$ meeting B does not meet $A \cap B$ and must therefore satisfy $N(C) \subseteq A \cap B \cap X_\gamma$, i.e. this component is not tight.

Given a tight component $C \in \mathcal{C}$ let $\mathcal{D}_C \subseteq \mathcal{D}$ be the set of all components D in \mathcal{D} with the property that there is some $(A, B) \in N_\gamma$ pointing away from X_γ such that D meets B and $C_B = C$. Informally, these sets \mathcal{D}_C are the components which we will need to group together with their C when choosing our next separations. The \mathcal{D}_C are pairwise disjoint: Indeed, given two separations (A, B) and (A', B') pointing away from X_γ , if $(B', A') \leq (A, B)$ then the set B' and B are disjoint, and if $(A, B) \leq (A', B')$, then $(A' \cap B') \setminus X_\gamma$ and $(A \cap B) \setminus X_\gamma$ cannot be contained in different tight components of $G - X_\gamma$.

Let $N_{\gamma+}$ consist of N_γ together with, for every tight component $C \in \mathcal{C}$ of $G - X_\gamma$, the separation $(C \cup \bigcup \mathcal{D}_C \cup X_\gamma, V(G) \setminus (C \cup \bigcup \mathcal{D}_C))$. It is easy to see that this set is a nested set of separations. Moreover, any pair of profiles efficiently distinguished by X_γ is efficiently distinguished by one of these new separations.

For limit ordinals γ let $N_\gamma := \bigcup_{\alpha < \gamma} N_\alpha$, this set is nested since every pair in N_γ is already in some N_α .

Then $N := N_\beta$ is the desired nested set of separations.

6.2.2. Canonical trees of tree-decompositions

To canonically and efficiently distinguish a robust set of principal profiles in a graph Carmesin, Hamann and Miraftab [3] introduced more complex objects than nested sets of separations: trees of tree-decompositions. These consist of a rooted tree where every node is associated with a tree-decomposition. At the root this is a tree decomposition of G . At every remaining node there is a tree-decomposition of one of the torsos of the tree-decomposition at the parent node. Their main result is the following:

THEOREM 1.5 ([3, Remark 8.3]) *Let G be a connected graph and \mathcal{P} a distinguishable set of principal robust profiles in G . There exists a canonical tree of tree-decompositions with the following properties:*

- (1) *the tree of tree-decompositions distinguishes \mathcal{P} efficiently;*
- (2) *if $t \in V(T)$ has level k , then (T_t, \mathcal{V}_t) contains only separations of order k ;*
- (3) *nodes t at all levels have $|V(T_t)|$ neighbours on the next level and the graphs assigned to them are all torsos of (T_t, \mathcal{V}_t) .*

We can also construct such a tree of tree-decompositions from our nested set of separators. In order to do that, let us recall the most important definitions from [3].

In a rooted tree (T, r) , the *level* of a vertex $t \in V(T)$ is $d(t, r) + 1$. A *tree of tree-decompositions* is a triple $((T, r), (G_t)_{t \in V(T)}, (T_t, \mathcal{V}_t)_{t \in V(T)})$ consisting of a rooted tree (T, r) , a family $(G_t)_{t \in V(T)}$ of graphs and a family $(T_t, \mathcal{V}_t)_{t \in V(T)}$ of tree-decompositions of the G_t . The graphs $G_{t'}$ assigned to the neighbours t' on the next level from a node $t \in V(T)$ shall be distinct torsos of the tree-decomposition (T_t, \mathcal{V}_t) . This tree of tree-decompositions is a tree of tree-decompositions of G , if $G_r = G$.

A separation (A, B) of G induces a separation (A', B') of G_t if $A \cap G_t = A'$ and $B \cap G_t = B'$. Given two profiles P, P' , we say that a tree of tree-decompositions (efficiently) distinguishes P and P' if there is a separation (A, B) in G (efficiently) distinguishing them and a node $t \in V(T)$ such that the separation induced by (A, B) on G_t is one of the separation induced by the tree-decomposition (T_t, \mathcal{V}_t) of G_t .

In order to deduce Theorem 1.5 from Theorem 1.3 it is useful to observe that our set of separators is nested in an even stronger sense: We say that two separators X and Y are

strongly nested if there is a component C of $G - X$ such that $Y \subseteq C \cup N(C)$ and there is a component C' of $G - Y$ such that $X \subseteq C' \cup N(C')$. The separators from the nested set \mathcal{N} from Theorem 1.3 are strongly nested:

LEMMA 6.16 *If X and Y are a pair of nested separators each of which efficiently distinguishes some pair of robust principal profiles, then they are strongly nested.*

Proof. We show that there is a component C of $G - X$ such that $Y \subseteq C \cup N(C)$.

If $Y \subseteq X$ the statement is obvious, by picking as C a tight component of $G - X$. So we may assume that Y meets some component C of $G - X$ in a vertex $v \in Y \cap C$. By nestedness $Y \subseteq C \cup X$. Suppose for a contradiction that $Y \not\subseteq C \cup N(C)$, i.e., Y contains a vertex $w \in X \setminus N(C)$.

Since Y efficiently distinguishes two principal profiles there are two distinct tight components C_1, C_2 of $G - Y$, by Lemma 6.1. X meets at most one of C_1 and C_2 since it is nested with Y ; without loss of generality we may assume $X \cap C_2 = \emptyset$. Since C_2 is a tight component of $G - Y$ there is a path P from v to w with all its interior vertices in C_2 . On the other hand v lies in C and w outside of $C \cup N(C)$, so $N(C)$ separates v from w . But $N(C) \subseteq X$ does not meet P since $X \cap C_2 = \emptyset$. This is a contradiction.

Note that for a separator X to be strongly nested with itself is a non-trivial property: It is precisely the statement that there is a tight component of $G - X$. Thus, if we talk about a strongly nested set of separators, we mean that not only any pair of distinct separators from that set is strongly nested, we also require each of the separators from that set to be nested with itself.

Next we show that we can close our strongly nested set under taking subsets:

LEMMA 6.17 *Let \mathcal{N} be a strongly nested set of separators and let \mathcal{N}' be the set of all subsets of elements of \mathcal{N} . Then \mathcal{N}' is strongly nested as well.*

Proof. Let $X, Y \in \mathcal{N}$ and let $X' \subseteq X, Y' \subseteq Y$, possibly equal. Take C_X to be a component of $G - X$ for which $Y \subseteq C_X \cup N(C_X)$, then in particular $Y' \subseteq C_X \cup N(C_X)$. Since $X' \subseteq X$ there is some component $C_{X'} \supseteq C_X$ of $G - X'$, thus $Y' \subseteq C_{X'} \cup N(C_{X'})$.

By symmetry we also find a component $C_{Y'}$ such that $X' \subseteq C_{Y'} \cup N(C_{Y'})$

So let \mathcal{N}' be the strongly nested set of all subsets of separators from \mathcal{N} , the canonical nested set of separators from Theorem 1.3. As such, \mathcal{N}' is canonical as well. The following lemma about separations with strongly nested separators will allow us to construct a tree of tree-decompositions from \mathcal{N}' inductively, starting with the separators of lowest size.

LEMMA 6.18 *If X, Y are distinct strongly nested separators and (A_X, B_X) and (A_Y, B_Y) are separations with separators X and Y respectively, such that $Y \subseteq B_X, X \subseteq B_Y$, then either (A_X, B_X) and (A_Y, B_Y) are nested, or there is a component C of $G - (X \cap Y)$ which meets neither X nor Y .*

Proof. Suppose that $(A_Y, B_Y) \not\subseteq (B_X, A_X)$. Then either there is a vertex in A_Y which does not lie in B_X , or there is a vertex in A_X which does not lie in B_Y . Since $A_X \cap B_X = X \subseteq B_Y$ and $A_Y \cap B_Y = Y \subseteq B_X$ either of these cases implies that there is a vertex v in $(A_X \setminus B_X) \cap (A_Y \setminus B_Y)$. This vertex v needs to lie in some component C of $G - (X \cup Y)$. However, C cannot send an edge to $X \setminus Y$ since such an edge would contradict the fact that (A_Y, B_Y)

is a separation. Similarly, C cannot be adjacent to any vertex of $Y \setminus X$. Thus C is in fact a component of $G - (X \cap Y)$ which meets either X nor Y .

Now we are ready to deduce Theorem 1.5 from Theorem 1.3:

Proof of Theorem 1.5. Let \mathcal{N}' be as above. We will build our tree of tree-decompositions inductively level-by-level, adding at stage k to every node t on level $k - 1$ new neighbours on level k , one for every torso of the tree-decompositions (T_t, \mathcal{V}_t) . We do this in a way that ensures the following properties:

- (i) if $d(r, t) = k$ then every separation in (T_t, \mathcal{V}_t) has order $k + 1$;
- (ii) every separator in \mathcal{N}' of size at least $k + 2$ is contained in exactly one of the torsos of (T_t, \mathcal{V}_t) whenever $d(t, r) \leq k$;
- (iii) if $d(t, r) = k$, every torso of (T_t, \mathcal{V}_t) meets at most one component of $G - X$ for every $X \in \mathcal{N}'$ of size $\leq k$ with $X \subseteq V(G_t)$.

Our inductive construction goes as follows: For $k = 0$ we consider the set S_1 which consists of, for every separator X of size 1 in \mathcal{N}' and every component C of $G - X$, the separation $(C \cup X, V(G) \setminus C)$, unless C is the only component of $G - X$.

Observe that S_1 is a nested set of separations: any two separations with the same separator are nested by construction and for separations with distinct separators X and Y the separators are disjoint, so $G - (X \cap Y) = G$ is connected and Lemma 6.18 gives that the separations are nested.

Moreover every ω -chain $(A_1, B_1) < (A_2, B_2) < \dots$ in S_1 has $\bigcap_{i \in \mathbb{N}} B_i = \emptyset$: We may assume without loss of generality that no two of these separations have the same separator since S_1 has no 3-chain of separations with the same separator. On the other hand a path from a vertex in $\bigcap_{i \in \mathbb{N}} B_i$ to A_1 (which has finite length) would need to meet all the infinitely many disjoint separators $A_i \cap B_i$.

Since S_1 contains no separation with a small orientation, it is, by construction, a regular tree set. Thus by Lemma 2.7 it induces a canonical tree-decomposition (T_r, \mathcal{V}_r) of $G_r = G$. We assign this tree-decomposition to the root of our tree of tree-decompositions and shall now verify (i) to (iii).

Observe that this decomposition satisfies (i) and (iii) as we only used separators of size 1 and every torso of (T_t, \mathcal{V}_t) meets at most one component of $G - X$ for every $X \in \mathcal{N}'$ of size ≤ 1 with $X \subseteq V(G_t)$. Moreover, (ii) is also satisfied since every separator X in \mathcal{N}' of size at least 2 is nested with each of the separators used in (T_r, \mathcal{V}_r) : such a separator cannot be contained in two distinct torsos since then a separation with separator in \mathcal{N}' would separate them. Conversely, there is a torso which contains X : otherwise consider a torso V_t that contains as much of X as possible and another torso V'_t which contains a vertex in $X \setminus V_t$. Then one of the edges on the path between t and t' in T again corresponds to a separation which separates X . But this is not possible since the separators of these separations are in \mathcal{N}' and thus nested with X .

For the k -th step of our construction, for $k \geq 1$, we attach at every node t on level $k - 1$ of our so-far constructed tree of tree-decompositions, for every torso G^t of (T_t, \mathcal{V}_t) a new node t' (which then is at level k) with $G_{t'} := G^t$. We then independently construct tree-decompositions for each of these torsos $G_{t'}$. For every torso we use all those separators from \mathcal{N}' which are of size $k + 1$ and lie inside that torso. Note that (ii) guarantees that every separator in \mathcal{N}' of size $k + 1$ is contained in exactly one of the newly added torsos.

Given one torso $G_{I'}$ of the tree-decomposition (T_t, \mathcal{V}_t) , we let S_{k+1} be the set of all separations (A, B) of $G_{I'}$ of order $k + 1$ with separator in \mathcal{N}' and the property that $A \setminus B$ is a component of $G - (A \cap B)$ but not the only one.

We claim that S_{k+1} is a nested set of separations. Indeed, if two separations from S_{k+1} with different separators X and Y were to cross then, by Lemma 6.18, there would be a component of $G_{I'} - (X \cap Y)$ avoiding X and Y . However, $X \cap Y$ has size less than k , lies in \mathcal{N}' and $G_{I'}$ meets, by (iii), at most one component of $G - (X \cap Y)$. Hence if we take vertices x and y in $G_{I'} - (X \cap Y)$ we find a path P between them in $G - (X \cap Y)$. But since $G_{I'}$ is obtained from G by repeatedly building a torso, $P \cap G_{I'}$ needs to contain a path between x and y in $G_{I'}$. In particular, this path does not meet $X \cap Y$ and thus $G_{I'} - (X \cap Y)$ has only one component, in particular every component of $G_{I'} - (X \cap Y)$ meets X and Y .

Now consider an ω -chain $(A_1, B_1) < (A_2, B_2) < \dots$ in S_{k+1} . We may assume without loss of generality that no two of these separations have the same separator, as in the case $k = 0$. If $\bigcap_{i \in \mathbb{N}} B_i$ is non-empty then its neighbourhood $Z := N_{G_{I'}}(\bigcap_{i \in \mathbb{N}} B_i)$ needs to be properly contained in some $A_l \cap B_l$: Every vertex in Z needs to be contained in some $A_m \cap B_m$ and if such a vertex lies in $A_m \cap B_m$, then it also lies in $A_n \cap B_n$ for every $n \geq m$. In particular, if $|Z| \geq k + 1$, there would be an m such that $A_m \cap B_m \subseteq Z$ and thus $A_n \cap B_n = A_m \cap B_m \forall n \geq m$ contradicting the assumption that no two of the (A_l, B_l) have the same separator. Hence $|Z| \leq k$ and we can easily find an l such that $Z \subsetneq A_l \cap B_l$.

But then again $G_{I'}$ would meet two distinct components of $G - Z$: one meeting $\bigcap_{i \in \mathbb{N}} B_i$ and one meeting A_l . This however is not possible since $|Z| < |A_l \cap B_l|$ and $Z \in \mathcal{N}'$.

By construction, S_{k+1} contains no separation with a small separation, it is thus a regular tree set. So, by Lemma 2.7, the set S_{k+1} induces a canonical tree-decomposition $(T_{I'}, \mathcal{V}_{I'})$ of $G_{I'}$. In this way we construct all the tree-decompositions for nodes at level k . We need to verify (i) to (iii). (i) is obvious. For (ii) we observe that every separator in \mathcal{N}' of size at least $k + 2$ which was contained in $G_{I'}$ was nested with every separator of a separation in S_{k+1} and is therefore contained in exactly one of the torsos of $(T_{I'}, \mathcal{V}_{I'})$, by the same argument as in the case $k = 0$.

For (iii) we note that for separators X of size $\leq k$ every torso of $(T_{I'}, \mathcal{V}_{I'})$ meets at most one component of $G - X$ as, by induction $G_{I'}$ itself only meets one component of $G - X$. For a separator X of size $k + 1$ let H be a torso of $(T_{I'}, \mathcal{V}_{I'})$. Firstly, H meets at most one component of $G_{I'} - X$ since if $G_{I'} - X$ has more than one component then X is one of the separators of $(T_{I'}, \mathcal{V}_{I'})$ and therefore, as S_{k+1} includes every separation of the form $(C \cup X, G_{I'} \setminus X)$ for any component C of $G_{I'} - X$, there needs to be a component C of $G_{I'} - X$ such that H is contained in $C \cup X$.

Secondly, when building the torso $G_{I'}$ from G we never add edges between distinct components of $G - X$ since we only add edges inside of separators in \mathcal{N}' , which are nested with X . Hence, if H would meet two components of $G - X$ it would also meet two components of $G_{I'} - X$. Hence H meets at most one component of $G - X$. This gives (iii).

Verification of correctness. Let us now verify that the so constructed tree of tree-decompositions $((T, r), (G_t)_{t \in V(t)}, (T_t, \mathcal{V}_t)_{t \in V(T)})$ – which is canonical by construction – has the properties (1) to (3) from the assertion. The properties (2) and (3) are fulfilled by construction, so we only need to verify (1).

Let P, P' be two robust principal profiles from \mathcal{P} . By Theorem 1.3, \mathcal{N}' contains some separator X which belongs to a separation efficiently distinguishing P and P' , say $|X| = k$. By our inductive construction, there is a unique G_t at level k which contains X . As P

and P' are principal profiles, there are two distinct components C, C' of $G - X$ such that $(V(G) \setminus C, C \cup X) \in P$, and $(V(G) \setminus C', C' \cup X) \in P'$. We claim that $C \cap V(G_t)$ is not empty.

Note that G_t is obtained from G by repeatedly taking some separation (A, B) of order $< k$ with $X \subseteq B$, deleting $A \setminus B$ and making $A \cap B$ complete. If we apply this operation for a single (A, B) which, say, turns some graph H with $V(H) \subseteq V(G)$ into H' then this preserves for H' the properties of H that (i) $H[C \cap V(H)]$ is connected and (ii) every vertex in X has, in H , a neighbour in $C \cap V(H)$. Thus every vertex in X has, in G_t , a neighbour in $C \cap V(G_t)$ proving that $C \cap V(G_t)$ is non-empty.

By a symmetrical argument not only C but also C' meets some component of $G_t - X$. Moreover, no two distinct components of $G - X$ can meet the same component of $G_t - X$: This would require an edge between these components, which would have to be added by the torso operation – but this operation only adds edges inside a separator $Y \in \mathcal{N}'$. And since Y is nested with X , that is Y meets only one component of $G - X$, this cannot add edges between different components of $G - X$.

Thus there is exactly one component C_t of $G_t - X$ such that $C_t \subseteq C$ and this component is not the only one from $G_t - X$. So, by construction the separation $(C_t \cup X, G_t \setminus X)$, which efficiently distinguishes the induced profiles of P and P' onto G_t is induced by (T_t, \mathcal{V}_t) .

REFERENCES

- [1] J. CARMESIN. All graphs have tree-decompositions displaying their topological ends. *Combinatorica* **39**(3) (2019), 545–596.
- [2] J. CARMESIN, R. DIESTEL, F. HUNDERTMARK, and M. STEIN. Connectivity and tree structure in finite graphs. *Combinatorica* **34**(1) (2014), 11–46.
- [3] J. CARMESIN, M. HAMANN and B. MIRAFTAB. Canonical trees of tree-decompositions. [arXiv:2002.12030](https://arxiv.org/abs/2002.12030) (2020).
- [4] R. DIESTEL. Ends and tangles. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* **87**(2) (2017), 223–244.
- [5] R. DIESTEL. *Graph Theory* Graduate Texts in Math. vol. **173** (Springer, Berlin, Heidelberg, 5th edition, 2017).
- [6] R. DIESTEL. Abstract separation systems. *Order* **35** (2018), 157–170.
- [7] R. DIESTEL, P. EBERENZ and J. ERDE. Duality theorems for blocks and tangles in graphs. *SIAM J. Discrete Math.* **31**(3) (2017), 1514–1528.
- [8] R. DIESTEL, J. ERDE and D. WEISSAUER. Structural submodularity and tangles in abstract separation systems. *J. Combin. Theory Ser. A*, **167C** (2019), 155–180.
- [9] R. DIESTEL, F. HUNDERTMARK and S. LEMANCZYK. Profiles of separations: in graphs, matroids, and beyond. *Combinatorica*, **39**(1) (2019), 37–75.
- [10] R. DIESTEL and J. KNEIP. Profinite separation systems. *Order*, **37**(1) (2020), 179–205.
- [11] R. DIESTEL and S. OUM. Tangle-tree duality in graphs, matroids and beyond. *Combinatorica* **39**(4) (2019), 879–910.
- [12] R. DIESTEL and S. OUM. Tangle-tree duality in abstract separation systems. *Adv. Math.* **377** (2021), 107470.
- [13] M. J. DUNWOODY and B. KRÖN. Vertex cuts. *J. Graph Theory* **80**(2) (2015), 136–171.
- [14] C. ELBRACHT, J. KNEIP and M. TEEGEN. Trees of tangles in abstract separation systems. *J. Combin. Theory Ser. A* **180** (2021), 105425.
- [15] C. ELBRACHT, J. KNEIP and M. TEEGEN. Obtaining trees of tangles from tangle-tree duality. To appear in *J. Comb.* [arXiv:2011.09758](https://arxiv.org/abs/2011.09758) (2020).
- [16] A.-K. ELM and J. KURKOFKA. A tree-of-tangles theorem for infinite tangles. [arXiv:2003.02535](https://arxiv.org/abs/2003.02535) (2020).

- [17] J. P. GOLLIN and J. KNEIP. Representations of infinite tree sets. *Order* **38**(1) (2021), 79–96.
- [18] R. HALIN. Lattices of cuts in graphs. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* **61** (1991), 217–230.
- [19] J. KNEIP. Ends as tangles. [arXiv:1909.12628](https://arxiv.org/abs/1909.12628) (2019).
- [20] N. ROBERTSON and P. D. SEYMOUR. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B* **52**(2) (1991), 153–190.