A tree-of-tangles theorem for infinite tangles

Ann-Kathrin Elm[1] · Jan Kurkofka[1]

Abstract

Carmesin has extended Robertson and Seymour's tree-of-tangles theorem to the infinite tangles of locally finite infinite graphs. We extend it further to the infinite tangles of all infinite graphs. Our result has a number of applications for the topology of infinite graphs, such as their end spaces and their compactifications.

1 Introduction

The *tree-of-tangles theorem*, one of the cornerstones of Robertson and Seymour's proof of their graph-minor theorem, says (in the terminology of [1, Sect. 12.5]):

Theorem *Every finite graph G has a nested set of separations which efficiently distinguishes all the finite tangles in G that can be distinguished.*

This is Theorem 12.5.4 in [1], the original article is [2].

Recently, Carmesin [3] has extended the tree-of-tangles theorem to the infinite tangles of infinite graphs that are locally finite. The precise statement of Carmesin's result reads:

Theorem *Every infinite connected graph G has a nested set of separations which efficiently distinguishes all the ends of G.*

✉ Jan Kurkofka
    jan.kurkofka@uni-hamburg.de

    Ann-Kathrin Elm
    ann-kathrin.elm@uni-hamburg.de

[1] Department of Mathematics, Universität Hamburg, Bundesstraße 55, 20146 Hamburg, Germany

Note that, in the wording of his theorem, Carmesin does not require the graph to be locally finite, and he speaks of ends where one expects infinite tangles. This is because his result is more general than an extension of the tree-of-tangles theorem to the infinite tangles of locally finite infinite graphs. To understand the difference, let us look at how the ends of a graph are related to its infinite tangles.

An end ω of a graph G (see [1]) orients every finite-order separation $\{A, B\}$ of G towards the side that contains a tail from every ray in ω. Since these orientations are, for distinct separations, consistent in a number of ways, they form an infinite tangle of G. Conversely, every infinite tangle of a locally finite and connected graph G is defined by an end in this way [4, 5]. Thus, if G is locally finite and connected, there is a canonical bijection between its infinite tangles and its ends. In this way, Carmesin's result extends the tree-of-tangles theorem to the infinite tangles of locally finite graphs.

When G is not locally finite, however, there can be infinite tangles that are not defined by an end. Then Carmesin's result no longer extends the tree-of-tangles theorem to the infinite tangles of G.

The infinite tangles that do not come from ends of the graph are fundamentally different from ends. They are closely related to free ultrafilters, and are called *ultrafilter tangles* [4]. More explicitly, by a recent result from [6], there is a canonical bijection between the ultrafilter tangles and the *ultrafilter tangle blueprints*: pairs (X, U) of a critical vertex set X and a free ultrafilter U on the collection $\mathscr{C}_X$ of the components C of $G - X$ for which the neighbourhood of C equals X, where a subset X of $V(G)$ is *critical* if X is finite and $\check{\mathscr{C}}_X$ is infinite. Therefore, every ultrafilter tangle $\tau = (X, U)$ has two aspects: Its combinatorial aspect is captured by its blueprint's critical vertex set X, and its ultrafilter aspect is encoded by the free ultrafilter U (see Sect. 2.7 for details). Since every vertex in a critical vertex set has infinite degree, it follows that locally finite connected graphs have no ultrafilter tangles, so all their infinite tangles are ends.

Ultrafilter tangles are interesting also for topological reasons. Every locally finite connected graph can be naturally compactified by its ends to form its well known end compactification [1] introduced by Freudenthal [7]. But for a non-locally finite graph, adding its ends no longer suffices to compactify it. Adding its ends *plus its ultrafilter tangles*, however, (i.e. adding all its infinite tangles) does again compactify the graph. This is Diestel's *tangle compactification* [4]. The tangle compactification generalises the end compactification twofold. On the one hand, it defaults to the end compactification when the graph is locally finite and connected. And on the other hand, the relation between the end compactification of locally finite connected graphs and their Stone-Čech compactification extends to all graphs when ends are generalised to tangles [8].

As our main result, we extend Robertson and Seymour's tree-of-tangles theorem to the infinite tangles of infinite graphs (and thus, we extend Carmesin's result from ends to all infinite tangles):

Theorem 1 *Every infinite connected graph G has a nested set of finite-order separations that efficiently distinguishes all the combinatorially distinguishable infinite tangles of G.*

Here, two infinite tangles are *combinatorially distinguishable* if not both are ultrafilter tangles or if both are ultrafilter tangles such that their critical vertex sets are distinct. So two infinite tangles are combinatorially indistinguishable if they are ultrafilter tangles with the same critical vertex. The proof will show that, conversely, the nested set of separations we find does not distinguish any two infinite tangles that are combinatorially indistinguishable.

As we will show, our result is best possible in the following sense. If a graph G has an ultrafilter tangle τ, then no nested set of finite-order separations of G efficiently distinguishes all the ultrafilter tangles of G that are not combinatorially distinguishable from τ.

1.1 Applications

Our work has four applications.

Elbracht, Kneip and Teegen use our Corollary 6.1 in their paper [9] to give a relatively short proof of Carmesin's tree-of-tangles theorem.

Call a set U of vertices in a graph G *tough* in G if only finitely many components of $G - X$ meet U for every finite vertex set $X \subseteq V(G)$. Bürger and the second author use our concept of 'corridors' in combination with our results from Sect. 5 to characterise the connected graphs G in which a given vertex set $U \subseteq V(G)$ is tough [10].

Our third application is the following structural connectivity result for infinite graphs, which generalises the way in which the cutvertices of a graph decompose it into its blocks in a tree-like fashion. Call a graph G *tough* if its vertex set is tough in G, i.e., if deleting finitely many vertices from G never leaves more than finitely many components. By the pigeonhole principle a graph is tough if and only if it has no critical vertex set.

Theorem 2 *Every connected graph G has a nested set of separations whose separators are precisely the critical vertex sets of G and all whose torsos are tough.*

(See Section 2.6 for definitions.)

Theorem 2 is interesting also from the perspective of topological infinite graph theory, in view of the following two results. Diestel and Kühn [11] showed that a graph is compactified by its ends if and only if it is tough (i.e., if and only if it has no critical vertex sets), and in [6] it was shown that every graph is compactified by its ends plus critical vertex sets. So a graph is compactified by points that come in two types, ends and critical vertex sets, and the second type decomposes the graph into a nested set of separations all whose torsos are compactified by the points of the first type.

Our fourth application answers a question that arises from the work of Polat and of Sprüssel. End spaces of graphs, in general, are not compact. However, Polat [12] and Sprüssel [13] independently showed that end spaces of graphs are normal. Polat even showed that end spaces of graphs are collectionwise normal, which is stronger than normal but weaker than compact Hausdorff. (In a *collectionwise* normal space one can at once pairwise separate any *collection* of closed disjoint sets with disjoint open neighbourhoods, cf. Definition 6.2.)

The infinite tangle space, endowed with the subspace topology of the tangle compactification, contains the end space as a subspace. As Diestel [4] showed, the infinite tangle space is compact Hausdorff, which implies collectionwise normality by general topology.

The ultrafilter tangle space, endowed with the subspace topology of the infinite tangle space, is not usually compact. Since the infinite tangle space is the disjoint union of the end space and the ultrafilter tangle space, the question arises whether the ultrafilter tangle space is collectionwise normal as well. We answer this question in the affirmative:

Theorem 3 *The ultrafilter tangle space of a graph is collectionwise normal.*

Our paper is organised as follows. Background knowledge is supplied in Sect. 2. In Sect. 3 we study examples and show that our main result is best possible. In Sect. 4 we give an overview on our overall proof strategy. Our main technical results are stated and proved in

Sect. 5. In Sect. 6 we provide the applications of our main technical results. In Sect. 7 we introduce an equivalence relation on a tree set given a consistent orientation of that tree set. This is the foundation for the definition of the modified torsos and proxies as well as for a 'lifting' process that we need in Sect. 8. In Sect. 8, finally, we introduce the modified torsos and prove our main result. Section 9 is our appendix.

Throughout this paper, $G = (V, E)$ is a connected graph of arbitrary cardinality.

2 Tools and terminology

We use the notation of Diestel's book [1]. For a short reminder on compactifications and inverse limits see the appendix Sect. 9.

2.1 Ends of graphs, and inverse limits

We write $\mathcal{X} = \mathcal{X}(G)$ for the collection of all finite subsets of the vertex set V of G, partially ordered by inclusion. An *end* of G, as defined by Halin [14], is an equivalence class of rays of G, where a ray is a one-way infinite path. Here, two rays are said to be *equivalent* if for every $X \in \mathcal{X}$ both have a subray (also called *tail*) in the same component of $G - X$. So in particular every end ω of G chooses, for every $X \in \mathcal{X}$, a unique component $C(X, \omega)$ of $G - X$ in which every ray of ω has a tail. In this situation, the end ω is said to *live* in $C(X, \omega)$. The set of ends of a graph G is denoted by $\Omega = \Omega(G)$. If $\mathcal{C}$ is any collection of components of $G - X$ for some $X \in \mathcal{X}$, we write $\Omega(X, \mathcal{C})$ for the set of ends ω of G with $C(X, \omega) \in \mathcal{C}$. The sets $\Omega(X, \mathcal{C})$ form a basis for a topology on Ω.

Recall that a *comb* is the union of a ray R (the comb's *spine*) with infinitely many disjoint finite paths, possibly trivial, that have precisely their first vertex on R. The last vertices of those paths are the *teeth* of this comb. Let us say that an end ω of G is contained *in the closure* of M, where M is either a subgraph of G or a set of vertices of G, if for every $X \in \mathcal{X}$ the component $C(X, \omega)$ meets M. Equivalently, ω lies in the closure of M if and only if G contains a comb with all its teeth in M and with its spine in ω. See [10, 15–17] for more on combs.

Now we describe an inverse system giving the end space: We note that $\mathcal{X}$ is directed by inclusion, and for every $X \in \mathcal{X}$ we let $\mathcal{C}_X$ be the set of components of $G - X$. Then letting $c_{X',X} : \mathcal{C}_{X'} \to \mathcal{C}_X$ for $X' \supseteq X$ send each component of $G - X'$ to the unique component of $G - X$ including it turns the sets $\mathcal{C}_X$ into an inverse system $\{\mathcal{C}_X, c_{X',X}, \mathcal{X}\}$. Clearly, its inverse limit consists precisely of the *directions* of the graph: choice maps f assigning to every $X \in \mathcal{X}$ a component of $G - X$ such that $f(X') \subseteq f(X)$ whenever $X' \supseteq X$. Every end ω of G defines a unique direction f_ω by mapping every $X \in \mathcal{X}$ to $C(X, \omega)$. In 2010, Diestel and Kühn [5] showed that, conversely, every direction in fact comes from a unique end in this way:

Theorem 2.1 ([5, Theorem 2.2]) *Let G be any graph. Then the map $\omega \mapsto f_\omega$ is a homeomorphism between the ends of G and its directions, i.e. $\Omega \cong \varprojlim \mathcal{C}_X$.*

2.2 Separations of sets and abstract separation systems

Separation systems are a fundamental notion in graph minor theory. In this section, we briefly recall the definitions from [1, 18, 19] that we need, without detailed explanations: for these we refer to the citations.

Fig. 1 The separations (A, B), (C, D), (E, F) form a star [1]
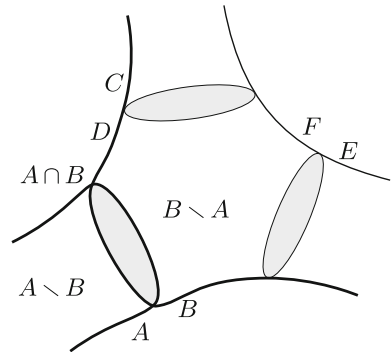
A *separation of a set* V is an unordered pair $\{A, B\}$ such that $A \cup B = V$. The ordered pairs (A, B) and (B, A) are its *orientations*. Then the *oriented separations* of V are the orientations of its separations. The map that sends every oriented separation (A, B) to its *inverse* (B, A) is an involution that reverses the partial ordering

$$(A, B) \leq (C, D) \;:\Leftrightarrow\; A \subseteq C \text{ and } B \supseteq D$$

since $(A, B) \leq (C, D)$ is equivalent to $(D, C) \leq (B, A)$.

More generally, a *separation system* is a triple $(\vec{S}, \leq, {}^{*})$ where $(\vec{S}, \leq)$ is a partially ordered set and ${}^{*}: \vec{S} \to \vec{S}$ is an order-reversing involution. We refer to the elements of $\vec{S}$ as *oriented separations*. If an oriented separation is denoted by $\vec{S}$, then we denote its *inverse* $\vec{S}^{*}$ as $\overleftarrow{s}$, and vice versa. That * is *order-reversing* means $\vec{r} \leq \vec{S} \leftrightarrow \overleftarrow{r} \geq \overleftarrow{s}$ for all $\vec{r}, \vec{S} \in \vec{S}$.

A *separation* is an unordered pair of the form $\{\vec{S}, \overleftarrow{s}\}$, and then denoted by s. Its elements $\vec{S}$ and $\overleftarrow{s}$ are the *orientations* of s. The set of all separations $\{\vec{S}, \overleftarrow{s}\} \subseteq \vec{S}$ is denoted by S. When a separation is introduced as s without specifying its elements first, we use $\vec{S}$ and $\overleftarrow{s}$ (arbitrarily) to refer to these elements. Every subset $S' \subseteq S$ defines a separation system $\vec{S'} := \bigcup S' \subseteq \vec{S}$ with the ordering and involution induced by $\vec{S}$.

Separations of sets, and their orientations, are an instance of this abstract setup if we identify $\{A, B\}$ with $\{(A, B), (B, A)\}$. Here is another example: The set $\vec{E}(T) := \{(x, y) \mid xy \in E(T)\}$ of all *orientations* (x, y) of the edges $xy = \{x, y\}$ of a tree T forms a separation system with the involution $(x, y) \mapsto (y, x)$ and the natural partial ordering on $\vec{E}(T)$ in which $(x, y) < (u, v)$ if and only if $xy \neq uv$ and the unique $\{x, y\}$–$\{u, v\}$ path in T links y to u.

In the context of a given separation system $(\vec{S}, \leq, {}^{*})$, a *star (of separations)* is a subset $\sigma \subseteq \vec{S}$ such that $\vec{r} \leq \overleftarrow{s}$ for all distinct $\vec{r}, \vec{S} \in \sigma$; see Fig. 1 for an illustration.[1] If t is a node of a tree T, then the set

---

[1] Officially, in [18] a star σ is additionally required to consist only of oriented separations $\vec{S}$ satisfying $\vec{S} \neq \overleftarrow{s}$. In this paper, however, all separations considered will satisfy this condition, which is why we will hide it for the convenience of the reader.

$$\vec{F}_t := \{ (x,t) \mid xt \in E(T) \}$$

is a star in $\vec{E}(T)$.

2.3 Orientations

An *orientation* of a separation system $\vec{S}$, or of a set S of separations, is a subset $O \subseteq \vec{S}$ such that $|O \cap \{\vec{s}, \overleftarrow{s}\}| = 1$ for all $s \in S$. A *partial orientation* of S is an orientation of a subset of S. A subset $O \subseteq \vec{S}$ is *consistent* if there are no two distinct separations $r, s \in S$ with orientations $\vec{r} < \vec{S}$ and $\overleftarrow{r}, \vec{S} \in O$. For example, the down-closure of any star $\vec{F}_t$ in $\vec{E}(T)$ is a consistent orientation.

2.4 Nested sets of separations and tree sets

Two separations are *nested* if they have comparable orientations. Two oriented separations $\vec{r}, \vec{S}$ are *nested* if r and s are nested. A set, either of separations or of oriented separations, is *nested* if every two of its elements are nested. For example, if T is a tree, then both $E(T)$ and $\vec{E}(T)$ are nested.

To state the definition of a tree set, we need the following definitions. An oriented separation $\vec{r} \in \vec{S}$ is

   (i)   *degenerate* if $\vec{r} = \overleftarrow{r}$,

   (ii)   *trivial* if there is a separation $s \in S$ such that both $\vec{r} < \vec{S}$ and $\vec{r} < \overleftarrow{s}$, and

   (iii)   *small* if $\vec{r} \le \overleftarrow{r}$.

In this paper, we will not have to worry about trivial separations. The only degenerate separation of a set V is (V, V); the small separations are precisely the ones of the form (A, V) with $A \subseteq V$. All degenerate and trivial separations are small.

A separation system is

   (i)   *essential* if it contains neither degenerate nor trivial elements, and

   (ii)   *regular* if it contains no small elements.

If $(\vec{S}, \le, *)$ is essential or regular, then we also call $\vec{S}$ and S *essential* or *regular*, respectively. Regular implies essential.

A *tree set* is a nested essential separation system. If $(\vec{S}, \le, *)$ is a tree set, then we also call $\vec{S}$ and S *tree sets*. If T is a tree, then $\vec{E}(T)$ is a tree set, the *edge tree set* of T.

In this paper, separations usually will not be small, and hence separation systems usually will be regular. This means that when we define a candidate for a tree set and have to verify that it really is a tree set, it will suffice to verify nestedness unless stated otherwise.

A consistent orientation O of a tree set $\vec{S}$ is equal to the down-closure

$$\lceil \sigma \rceil_{\vec{S}} := \{ \vec{r} \in \vec{S} \mid \exists \vec{S} \in \sigma : \vec{r} \le \vec{S} \}$$

in $\overrightarrow{S}$ of the star σ formed by the maximal elements of O if and only if every element of O lies below some maximal element of O. We call these orientation defining stars σ the *splitting stars* of $\overrightarrow{S}$. For example, the splitting stars of the edge tree set $\overrightarrow{E}(T)$ of a finite tree are precisely the stars $\overrightarrow{F}_t$. But if T is a ray $v_0 v_1 \ldots$, then $\{ (v_n, v_{n+1}) \mid n \in \mathbb{N} \}$ is a consistent orientation of $\overrightarrow{E}(T)$ that has no maximal element.

Gollin and Kneip [20] characterised the tree sets that are isomorphic to the edge tree set of a tree. An *isomorphism* between two separation systems is a bijection between their underlying sets that respects both their partial orderings and their involutions. A chain $\mathcal{C}$ in a given poset is said to have *order-type* α for an ordinal α if $\mathcal{C}$ with the induced linear order is order-isomorphic to α. The chain $\mathcal{C}$ is then said to be an α-*chain*.

Theorem 2.2 ([20, Theorem 1]) *A tree set is isomorphic to the edge tree set of a tree if and only if it is regular and contains no $(\omega + 1)$-chain.*

2.5 Separations of graphs

A *separation of a graph* G is a separation $\{A, B\}$ of the set $V(G)$ (meaning $A \cup B = V(G)$) such that G has no edge 'jumping' the *separator* $A \cap B$, meaning that G contains no edge between $A \setminus B$ and $B \setminus A$. Thus, (oriented) separations of graphs are an instance of (oriented) separations of sets. The *order* of $\{A, B\}$ is the cardinal $|A \cap B|$. The set of all finite-order separations of a graph G is denoted by $S_{\aleph_0} = S_{\aleph_0}(G)$. A tree set *of* G is a tree set of separations of G with the usual partial ordering and involution.

If (A, B) and (C, D) are two separations of G, then

(i) $(A, B) \vee (C, D) := (A \cup C, B \cap D)$ is their *supremum*, and
(ii) $(A, B) \wedge (C, D) := (A \cap C, B \cup D)$ is their *infimum*.

Supremum and infimum satisfy De Morgan's law: $(\overrightarrow{r} \vee \overrightarrow{S})^* = \overleftarrow{r} \wedge \overleftarrow{s}$.

The following non-standard notation often will be useful as an alternative perspective on separations of graphs. Recall that for a vertex set $X \subseteq V(G)$ we denote the collection of the components of $G - X$ by $\mathcal{C}_X$. If any $X \subseteq V(G)$ and $\mathcal{C} \subseteq \mathcal{C}_X$ are given, then these give rise to a separation of G which we denote by

$$\{X, \mathcal{C}\} := \{ V \setminus V[\mathcal{C}] \,,\, X \cup V[\mathcal{C}] \}$$

where $V[\mathcal{C}] = \bigcup \{ V(C) \mid C \in \mathcal{C} \}$. Note that every separation $\{A, B\}$ of G with $A, B \subseteq V(G)$ can be written in this way. For the orientations of $\{X, \mathcal{C}\}$ we write

$$(X, \mathcal{C}) := (V \setminus V[\mathcal{C}] \,,\, X \cup V[\mathcal{C}]) \quad \text{and} \quad (\mathcal{C}, X) := (V[\mathcal{C}] \cup X \,,\, V \setminus V[\mathcal{C}]).$$

If C is a component of $G - X$ we write $\{X, C\}$ instead of $\{X, \{C\}\}$. Similarly, we write (C, X) and (X, C) instead of $(\{C\}, X)$ and $(X, \{C\})$, respectively.

2.6 Parts and torsos

If T is a tree set of separations of G and O is a consistent orientation of T, then the intersection $\Pi = \bigcap \{ B \mid (A, B) \in O \}$ is called the *part* of O. And the graph that is obtained from $G[\Pi]$ by adding an edge xy whenever $x \neq y \in \Pi$ lie together in the separator of some

separation of O is called the *torso* of O (or of Π if O is clear from context). We denote the torso of O by $\mathrm{torso}(G, O)$.

We will need the following lemma and its corollaries (the lemma is folklore and has been proved, e.g., in [3]; we present an alternative proof for convenience):

Lemma 2.3 *If Π is a part of a tree set of G, then for every $G[\Pi]$-path P there is some separation of the tree set whose separator contains both endvertices of P.*

Proof Let O be any consistent orientation of a tree set of G, write Π for its part and suppose that $P = xv_1 \ldots v_n y$ is a $G[\Pi]$-path (so $n \geq 1$). For every $k \in [n]$ pick an oriented separation $(A_k, B_k) \in O$ with $v_k \in A_k \setminus B_k$ (so that (A_k, B_k) witnesses $v_k \notin \Pi$). Let N consist of the $\leq$-maximal separations from the collection $\{ (A_k, B_k) \mid k \in [n] \}$. Then for every v_k there is a separation $(A, B) \in N$ with $v_k \in A \setminus B$. Our aim is to show that N is a singleton, since then the separator of the sole separation in N must contain both x and y, so we would be done. By the choice of N, every two oriented separations in N are $\leq$-incomparable. As O is a consistent orientation of a tree set, this means that N must be a star. Then $|N| = 1$ is evident, since otherwise the sides $G[A \setminus B]$ for $(A, B) \in N$ altogether induce a disconnection of the subpath $v_1 \ldots v_n$ of P contradicting its connectedness. $\qquad\square$

Corollary 2.4 *If Π is a part of a tree set of G and ω is an end of G in the closure of Π while $G[\Pi]$ coincides with the torso of Π, then ω has a ray in $G[\Pi]$.*

Proof If ω lies in the closure of Π, we find a comb in G with its spine R in ω and all of its teeth in Π. Without loss of generality the comb meets Π precisely in its teeth. Then, as $G[\Pi]$ coincides with the torso of Π, it has an edge between every two consecutive teeth by Lemma 2.3, and so contains a ray equivalent to R. $\qquad\square$

Corollary 2.5 *If Π is a part of a tree set of G and two rays of $G[\Pi]$ are equivalent in G, then they are equivalent in the torso of Π as well.*

Proof Given two rays of $G[\Pi]$ that are equivalent in G, we inductively construct infinitely many pairwise vertex-disjoint paths in G between them, and then employ Lemma 2.3 to turn these into paths of the torso. $\qquad\square$

The next corollary has already been known to Carmesin [3]:

Corollary 2.6 *The intersection of a connected set of vertices of G with a part of a tree set of separations of G induces a connected subgraph of the part's torso.* $\qquad\square$

2.7 Infinite tangles

The *interior* of a star $\{ (A_i, B_i) \mid i \in I \} \subseteq \vec{S}_{\aleph_0}$ is the intersection $\bigcap_{i \in I} B_i$. An $\aleph_0$-*tangle* (of G), or *infinite tangle*, is a consistent orientation of $S_{\aleph_0}$ that contains no finite star of finite interior as a subset. We write $\Theta = \Theta(G)$ for the set of all $\aleph_0$-tangles of G. This particular definition is due to Diestel who showed in his paper [4] that it is equivalent to the original

definition by Robertson and Seymour [2]. Infinite tangles are resilient in the following sense:

Lemma 2.7 ([4, *Lemma* 1.10]) *Let τ be an $\aleph_0$-tangle of G and $(A, B) \in \tau$. Let (A', B') be a separation of G with $A \triangle A'$ and $B \triangle B'$ finite. Then $(A', B') \in \tau$.*

In the remainder of this section, we give a summary of the results in Diestel's paper [4] on infinite tangles. If ω is an end of G, then letting

$$\tau_\omega := \big\{ (X, \mathscr{C}) \in \vec{S}_{\aleph_0} \mid C(X, \omega) \in \mathscr{C} \big\}$$

defines an injection $\Omega \hookrightarrow \Theta$, $\omega \mapsto \tau_\omega$. The $\aleph_0$-tangles of the form τ_ω are called *end tangles*. By abuse of notation we write Ω for the collection of all end tangles of G, so we have $\Omega \subseteq \Theta$.

In order to understand the $\aleph_0$-tangles that are not ends, Diestel studied an inverse limit description of Θ. In this paper, partition classes are required to be non-empty as usual, with the exception that whenever we speak of a *bipartition* we allow for at most one empty class. Now if τ is an $\aleph_0$-tangle, then for every $X \in \mathcal{X}$ it chooses one *big side* from each bipartition $\{\mathscr{C}, \mathscr{C}'\}$ of $\mathscr{C}_X$, namely the $\mathscr{D} \in \{\mathscr{C}, \mathscr{C}'\}$ with $(X, \mathscr{D}) \in \tau$. Since it chooses theses sides consistently, it induces an ultrafilter $U(\tau, X)$ on $\mathscr{C}_X$, one for every $X \in \mathcal{X}$, which is given by

$$U(\tau, X) = \{ \mathscr{C} \subseteq \mathscr{C}_X \mid (X, \mathscr{C}) \in \tau \},$$

and these ultrafilters are compatible in that they form a limit of the inverse system $\{ \beta(\mathscr{C}_X), \beta(c_{X'X}), \mathcal{X} \}$. Here, each set $\mathscr{C}_X$ is endowed with the discrete topology and $\beta(\mathscr{C}_X)$ denotes its Stone-Čech compactification. Every bonding map $\beta(c_{X'X})$ is the unique continuous extension of $c_{X'X}$ that is provided by the Stone-Čech property. More explicitly, the map $\beta(c_{X'X})$ sends each ultrafilter $U' \in \beta(\mathscr{C}_{X'})$ to its *restriction*

$$U' {\restriction} X = \{ \mathscr{C} \subseteq \mathscr{C}_X \mid \exists \mathscr{C}' \in U' : \mathscr{C} \supseteq \mathscr{C}' {\restriction} X \} \in \beta(\mathscr{C}_X)$$

where $\mathscr{C}' {\restriction} X = c_{X'X}[\mathscr{C}']$. As one of his main results, Diestel showed that the map

$$\tau \mapsto (U(\tau, X) \mid X \in \mathcal{X})$$

defines a bijection between the tangle space Θ and the inverse limit $\varprojlim \beta(\mathscr{C}_X)$. Moreover, he showed that the ends of G are precisely those $\aleph_0$-tangles whose induced ultrafilters are all principal.

For every $\aleph_0$-tangle τ we write $\mathcal{X}_\tau$ for the collection of all $X \in \mathcal{X}$ for which the induced ultrafilter $U(\tau, X)$ is free. Equivalently, $\mathcal{X}_\tau$ is the collection of those $X \in \mathcal{X}$ for which the star $\{ (C, X) \mid C \in \mathscr{C}_X \}$ is included in τ. The set $\mathcal{X}_\tau$ is empty if and only if τ is an end tangle. An $\aleph_0$-tangle τ with $\mathcal{X}_\tau$ non-empty is called an *ultrafilter tangle*, and we write Υ for the collection of all ultrafilter tangles, i.e. $\Upsilon = \Theta \setminus \Omega$.

Theorem 2.8 ([4, *Theorem* 3.5]) *For every ultrafilter tangle τ and each $X \in \mathcal{X}_\tau$ the free ultrafilter $U(\tau, X)$ determines τ in that*

$$\tau = \big\{ (A, B) \in \vec{S}_{\aleph_0} \mid \exists \mathscr{C} \in U(\tau, X) : V[\mathscr{C}] \subseteq B \setminus A \big\}.$$

For every ultrafilter tangle τ the set $\mathcal{X}_\tau \subseteq \mathcal{X}$ has a least element X_τ of which it is the up-closure $\mathcal{X}_\tau = \lfloor X_\tau \rfloor_{\mathcal{X}} := \{ X \in \mathcal{X} \mid X \supseteq X_\tau \}$. These elements have been characterised combinatorially in [6] as follows. Given $X \subseteq V(G)$ and a subset $Y \subseteq X$ we write $\mathscr{C}_X(Y)$ for the set $\{ C \in \mathscr{C}_X \mid N(C) = Y \}$ of components of $G - X$ that have their neighbourhood precisely equal to Y. In the special case of $X = Y$ we abbreviate $\mathscr{C}_X(X)$ to $\breve{\mathscr{C}}_X$. A finite vertex set $X \in \mathcal{X}$ is *critical* if $\breve{\mathscr{C}}_X$ is infinite. The collection of all the critical vertex sets of a graph G is denoted by $\mathrm{crit}(G)$. A vertex set $X \subseteq V(G)$ is of the form X_τ for an ultrafilter tangle τ of G if and only if X is critical. But more is true: The critical vertex sets allow us to describe the ultrafilter tangles explicitly, like the ends allow us to describe the end tangles explicitly. An *ultrafilter tangle blueprint* is an ordered pair (X, U) of a critical vertex set X and a free ultrafilter U on $\breve{\mathscr{C}}_X$.

Theorem 2.9 ([6, *Theorem* 4.10]) *Let G be any graph. Then the map*

$$\tau \mapsto \left(X_\tau, \, U(\tau, X_\tau) \cap 2^{\breve{\mathscr{C}}_{X_\tau}} \right)$$

is a bijection between the ultrafilter tangles and the ultrafilter tangle blueprints. In particular, $(X_\tau, \breve{\mathscr{C}}_{X_\tau}) \in \tau$ for every ultrafilter tangle τ of G.

We will resume the following notation from [6] for critical vertex sets. For every $X \in \mathcal{X}$ and all critical Y that are not entirely contained in X we write $C_X(Y)$ for the unique component of $G - X$ meeting Y (equivalently: including $\bigcup \mathscr{C}_{X \cup Y}(Y)$).

Lemma 2.10 ([6, *Lemma* 4.8]) *For every ultrafilter tangle τ and each $X \in \mathcal{X} \setminus \mathcal{X}_\tau$ we do have $X_\tau \subseteq X \cup C_X(X_\tau)$ and the ultrafilter $U(\tau, X)$ is generated by $\{ C_X(X_\tau) \}$.*

Next, we describe Diestel's tangle compactification. For this, we recall the notion of the 1-complex of a graph G. In the *1-complex* of G which we denote also by G, every edge $e = xy$ is a homeomorphic copy $[x, y] := \{x\} \sqcup \mathring{e} \sqcup \{y\}$ of $[0, 1]$ with $\mathring{e}$ corresponding to $(0, 1)$. The point set of G is $V \sqcup \bigsqcup_{e \in E} \mathring{e}$. Points in $\mathring{e}$ are called *inner edge points*, and they inherit their basic open neighbourhoods from $(0, 1)$. For each subcollection $F \subseteq E$ we write $\mathring{F}$ for the set $\bigsqcup_{e \in F} \mathring{e}$ of inner edge points of edges in F. The basic open neighbourhoods of a vertex v of G are given by unions $\bigcup_{e \in E(v)} [v, i_e)$ of half open intervals with each i_e some inner edge point of e where $E(v)$ denotes the set of edges of G at v.

To obtain the tangle compactification $|G|_\Theta$ of a graph G we extend the 1-complex of G to a topological space $G \sqcup \Theta = G \sqcup \varprojlim \beta(\mathscr{C}_X)$ by declaring as open in addition to the open sets of G, for all $X \in \mathcal{X}$ and all $\mathscr{C} \subseteq \mathscr{C}_X$, the sets

$$\mathcal{O}_{|G|_\Theta}(X, \mathscr{C}) := \bigcup \mathscr{C} \cup \mathring{E}\left(X, \bigcup \mathscr{C}\right) \cup \left\{ \, (U_Y : Y \in \mathcal{X}) \in \varprojlim \beta(\mathscr{C}_X) \, \middle| \, \mathscr{C} \in U_X \, \right\}$$

and taking the topology this generates.

Theorem 2.11 ([4, *Theorem* 1]) *Let G be any graph.*

(i) $|G|_\Theta$ *is a compactification of G with totally disconnected remainder.*

(ii) *If G is locally finite and connected, then $|G|_\Theta$ coincides with the Freudenthal compactification of G.*

Theorem 2.12 ([8]) *The tangle compactification of any graph G is obtained from its Stone-Čech compactification βG by first declaring G to be open[2] in βG and then collapsing each connected component of the Stone-Čech remainder to a single point.*

2.8 Trees-of-tangles

A separation $\{A, B\}$ of G and its orientations *distinguish* two infinite tangles τ_1 and τ_2 of G if τ_1 and τ_2 orient $\{A, B\}$ differently, i.e., if $(A, B) \in \tau_1$ and $(B, A) \in \tau_2$ or vice versa. The separation $\{A, B\}$ distinguishes τ_1 and τ_2 *efficiently* if it has minimal order $|A \cap B|$ among all the separations of G that distinguish τ_1 and τ_2. The result by Carmesin that implies the tree-of-tangles theorem for the infinite tangles of locally finite infinite graphs states:

Theorem 2.13 ([3, *Corollary* 5.17]) *Every connected graph G has a tree set of finite-order separations of G that efficiently distinguishes all the ends of G.*

Here, we view the ends as end tangles.

Very recently, Carmesin, Hamann and Miraftab showed a canonical version of this theorem; see their paper [21]. We should also mention the work by Dunwoody and Krön [22] and the work by Elbracht, Kneip and Teegen [9].

2.9 Combinatorial indistinguishability and tame separations

We say that two infinite tangles τ_1 and τ_2 are *combinatorially distinguishable* if at least one of them is an end tangle or they are both ultrafilter tangles but such that their critical vertex sets X_{τ_1} and X_{τ_2} are distinct. Thus, when τ_1 and τ_2 are combinatorially indistinguishable, they are ultrafilter tangles with $X_{\tau_1} = X_{\tau_2}$. Then we also call them *equivalent* for short and write $\tau_1 \sim \tau_2$. There exist separations of G that do not distinguish any two equivalent tangles:

Definition 2.14 A finite-order separation $\{X, \mathscr{C}\}$ of G and its orientations are *tame* if for no $Y \subseteq X$ both $\mathscr{C}_X(Y) \cap \mathscr{C}$ and $\mathscr{C}_X(Y) \cap (\mathscr{C}_X \setminus \mathscr{C})$ are infinite. We write $S_t = S_t(G)$ for the set of all tame finite-order separations of G.

Lemma 2.15 ([6, §5]) *If τ_1 and τ_2 are two equivalent infinite tangles of a graph G, then $\tau_1 \cap \vec{S}_t = \tau_2 \cap \vec{S}_t$.*

The tree set that we construct in the proof of Theorem 1 will consist of tame separations.

Identifying all equivalent $\aleph_0$-tangles yields the quotient $\Theta/\sim = \Omega \sqcup \mathrm{crit}(G)$ which is yet again a tangle space. An $\aleph_0$-*tangle of* S_t is a consistent orientation of S_t that contains no finite star of finite interior as a subset. We write $\Theta_t = \Theta_t(G)$ for the set of all $\aleph_0$-tangles of $S_t(G)$.

Theorem 2.16 ([6, *Theorem* 5.10]) *Let G be any graph. The $\aleph_0$-tangles of S_t are precisely the ends and critical vertex sets of G, i.e. $\Theta_t(G) = \Omega(G) \sqcup \mathrm{crit}(G)$.*

---

[2] When G is locally compact, it is automatically open in βG, and so this step is redundant for locally finite graphs.

As a consequence, every graph is compactified by its ends and critical vertex sets in the tangle-type compactification $|G|_\Theta / \sim \, = G \sqcup \Omega(G) \sqcup \mathrm{crit}(G) = G \sqcup \Theta_t(G)$.

Since the tree set that we construct in the proof of Theorem 1 will efficiently distinguish all the combinatorially distinguishable $\aleph_0$-tangles of G and consist of tame separations, it will efficiently distinguish all the $\aleph_0$-tangles of S_t (i.e. all the ends and critical vertex sets of G).

3 Example section

The aim of this section is twofold. First, we verify that our main result, Theorem 1, is indeed best possible as claimed in the introduction. More precisely, in Sect. 3.1 we show that tree sets of finite-order separations cannot distinguish all the ultrafilter tangles from the same equivalence class at once—for any G.

Second, we study the candidate for a starting tree set that is formed by the separations $\{X, \breve{\mathscr{C}}_X\}$ with X critical in G (recall that these are precisely the separations which naturally accompany the ultrafilter tangles). More precisely, in Sect. 3.2 we will see two example graphs showing that it is necessary to modify the tree set candidate: For the first example graph, the separations $\{X, \breve{\mathscr{C}}_X\}$ form a tree set but do not distinguish any two ultrafilter tangles at all. For the second example graph, the separations $\{X, \breve{\mathscr{C}}_X\}$ are not even nested.

3.1 Ultrafilters and tree sets

In this subsection we show that, as soon as a graph G has some ultrafilter tangle τ, it already cannot admit a tree set of finite-order separations that distinguishes all the ultrafilter tangles that are equivalent to τ. As our first step, we translate the problem from graphs to bipartitions of sets.

For this, we need to make some things formal first. Suppose that K is a non-empty set. We let $\vec{\mathcal{B}}(K) := 2^K$. Thus, every subset of K is an oriented 'separation'. The partial ordering $\leq$ of $\vec{\mathcal{B}}(K)$ will be $\supseteq$, the involution * on $\vec{\mathcal{B}}(K)$ will be complementation in the set K. If desired, we can think of a separation $Z \subseteq K$ as the oriented bipartition (Z^*, Z) of K, and then $\mathcal{B}(K)$ is the set of bipartitions of K. Note that two separations $Z_1, Z_2 \in \vec{\mathcal{B}}(K)$ are nested if $Z_1 \subseteq Z_2$ or $Z_1 \supseteq Z_2$ or $Z_1 \cup Z_2 = K$ or $Z_1 \cap Z_2 = \emptyset$. A *tree set of bipartitions* of K is a tree set contained in $\vec{\mathcal{B}}(K)$ with the induced partial ordering and involution. Note that $\emptyset$ is the sole small separation in $\vec{\mathcal{B}}(K)$ for $Z \subseteq K \setminus Z$ implies $Z = \emptyset$. Since an ultrafilter on K happens to be an orientation of $\mathcal{B}(K)$, a tree set $\vec{T}$ of bipartitions of K distinguishes two distinct ultrafilters $U \neq U'$ on K if there is some $Z \in \vec{T}$ with $Z \in U$ and $Z^* \in U'$. We are almost ready for the translation, we only need one more lemma:

Lemma 3.1 *Let τ be any ultrafilter tangle of G with blueprint (X, U) and let any separation $(Y, \mathscr{D}) \in \tau$ be given. Write $\mathscr{C}$ for the set of those components in $\breve{\mathscr{C}}_X$ that avoid Y.*

  (i)   *If Y includes X, then $(Y, \mathscr{D}) \leq (X, \mathscr{D} \cap \breve{\mathscr{C}}_X) \in \tau$.*
  (ii) *Otherwise $(Y, \mathscr{D}) \leq (X, \mathscr{C}) \in \tau$.*

In particular, the set $\big\{ (X, \mathscr{C}) \mid \mathscr{C} \in U \big\}$ is cofinal in τ.

Proof Since Y is finite, $\mathscr{C}$ is a cofinite subset of $\breve{\mathscr{C}}_X$, giving $\mathscr{C} \in U$.

(i) The intersection $\mathscr{D} \cap \breve{\mathscr{C}}_X$ can be written in a more complicated way as $(\mathscr{D}{\restriction}X) \cap \mathscr{C}$ where $\mathscr{C} \in U$ as noted above and

$$\mathscr{D}{\restriction}X = \{\, C \in \mathscr{C}_X \mid \exists D \in \mathscr{D} : D \supseteq C \,\} \in U.$$

Hence $(X, \mathscr{D} \cap \breve{\mathscr{C}}_X) \in \tau$. It is straightforward to check $(Y, \mathscr{D}) \leq (X, \mathscr{D} \cap \breve{\mathscr{C}}_X)$.

(ii) From $\mathscr{C} \in U$ we get $(X, \mathscr{C}) \in \tau$. Lemma 2.10 deduces from $(Y, \mathscr{D}) \in \tau$ that $C_Y(X) \in \mathscr{D}$. Finally, we calculate $(Y, \mathscr{D}) \leq (Y, C_Y(X)) \leq (X, \mathscr{C})$ where for the second inequality we use that every component in $\mathscr{C}$ sends an edge to the non-empty $X \setminus Y \subseteq C_Y(X)$ to deduce $\bigcup \mathscr{C} \subseteq C_Y(X)$. $\qquad\square$

Now we are ready for the translation:

Lemma 3.2 *Let X be a critical vertex set of G. Then every tree set of finite-order separations of G that distinguishes all the ultrafilter tangles τ of G with $X_\tau = X$ does induce a tree set of bipartitions of $\breve{\mathscr{C}}_X$ that distinguishes all the free ultrafilters on $\breve{\mathscr{C}}_X$.*

Proof Let $\overrightarrow{T}$ be a tree set of finite order separations of G that distinguishes all the ultrafilter tangles of G with $X_\tau = X$. Without loss of generality every separation $(Y, \mathscr{D}) \in \overrightarrow{T}$ distinguishes some two such ultrafilter tangles, and so $X \subseteq Y$ follows for all $(Y, \mathscr{D}) \in \overrightarrow{T}$.

The candidate for a tree set of bipartitions of $\breve{\mathscr{C}}_X$ is $\{\, \bar{\mathscr{D}} \mid (Y, \mathscr{D}) \in \overrightarrow{T} \,\}$ where $\bar{\mathscr{D}} = \mathscr{D} \cap \breve{\mathscr{C}}_X$. But when $(Y, \mathscr{D}')$ is the inverse of $(Y, \mathscr{D})$ it can happen that $\bar{\mathscr{D}}'$ is not the inverse of $\bar{\mathscr{D}}$ in $\mathcal{B}(\breve{\mathscr{C}}_X)$. For example, this happens when a finite component $C \in \breve{\mathscr{C}}_X$ is contained in Y, for then both $\bar{\mathscr{D}}'$ and $\bar{\mathscr{D}}$ are missing C.

We overcome this obstacle as follows. First, we choose any consistent orientation O of $\overrightarrow{T}$ (such an orientation exists, e.g., by [18, Lemma 4.1] which essentially applies Zorn's lemma to achieve this). Then, we define $N_O := \{\, \bar{\mathscr{D}} \mid (\mathscr{D}, Y) \in O \,\}$. Finally, we claim that $\overrightarrow{N} := N_O \cup N_O^*$ is a tree set of bipartitions of $\breve{\mathscr{C}}_X$ that distinguishes all the free ultrafilters on $\breve{\mathscr{C}}_X$.

To verify that $\overrightarrow{N}$ is a tree set we show that $\overrightarrow{N}$ is nested. For this, consider any two separations $(\mathscr{D}_1, Y_1), (\mathscr{D}_2, Y_2) \in O$. Then, say, either $(\mathscr{D}_1, Y_1) \leq (\mathscr{D}_2, Y_2)$ implies $\bar{\mathscr{D}}_1 \subseteq \bar{\mathscr{D}}_2$ or $(\mathscr{D}_1, Y_1) \leq (Y_2, \mathscr{D}_2)$ implies $\bar{\mathscr{D}}_1 \subseteq (\bar{\mathscr{D}}_2)^*$. So $\overrightarrow{N}$ is a tree set.

Now let $U \neq U'$ be any distinct two free ultrafilters on $\breve{\mathscr{C}}_X$. Then there is a separation $(\mathscr{D}, Y) \in O$ that distinguishes the ultrafilter tangles τ_U and $\tau_{U'}$ corresponding to (X, U) and (X, U'), say with $(\mathscr{D}, Y) \in \tau_U$ and $(Y, \mathscr{D}) \in \tau_{U'}$. By Lemma 3.1 we have $\bar{\mathscr{D}} \in U'$. Similarly $\overline{\mathscr{D}_Y \setminus \mathscr{D}} \in U$, which then via the inclusion $\overline{\mathscr{D}_Y \setminus \mathscr{D}} \subseteq (\bar{\mathscr{D}})^*$ implies $(\bar{\mathscr{D}})^* \in U$. $\qquad\square$

As a consequence of this lemma, it suffices to show

Theorem 3.3 *If K is an infinite set, then no tree set of bipartitions of K distinguishes all the free ultrafilters on K.*

in order to obtain our desired result:

Corollary 3.4 *If τ is an ultrafilter tangle of G, then no tree set of finite-order separations of G distinguishes all the ultrafilter tangles that are equivalent to τ.* ☐

Theorem 3.3 above has been proved independently from us by Bowler [23] in 2014 who did not publish his findings. The proof presented below is ours. For the proof we need the following lemma which is a tree set version of the fact that every connected infinite graph contains either a ray or a vertex of infinite degree, [1, Proposition 8.2.1]:

Lemma 3.5 *Every regular infinite tree set contains either an ω-chain or an infinite splitting star.*

Proof If a tree set contains no ω-chain, then it is isomorphic to the edge tree set of a rayless tree by Theorem 2.2. This tree, then, must have an infinite degree vertex if the tree set is infinite. ☐

If U is an ultrafilter on a set K and $\mathcal{K}$ is a partition of K, then we write $U.\mathcal{K}$ for the *induced* ultrafilter on $\mathcal{K}$ given by $\{\mathcal{A} \subseteq \mathcal{K} \mid \bigcup \mathcal{A} \in U\}$. Notably, if U is principal, then so is $U.\mathcal{K}$. Conversely, every ultrafilter $\mathcal{U}$ on $\mathcal{K}$ gives a filter

$$\lfloor \{\bigcup \mathcal{A} \mid \mathcal{A} \in \mathcal{U}\} \rfloor_K := \{A \subseteq K \mid \exists \mathcal{A} \in \mathcal{U} : A \supseteq \bigcup \mathcal{A}\}$$

on K, and every ultrafilter U on K that extends this filter induces $\mathcal{U}$ in that $\mathcal{U} = U.\mathcal{K}$. Phrased differently, the map $U \mapsto U.\mathcal{K}$ is a surjection from the set of ultrafilters on K onto the set of ultrafilters on $\mathcal{K}$. Notably, free ultrafilters on $\mathcal{K}$ are induced only by free ultrafilters on K.

Proof of Theorem 3.3 Let any infinite set K be given and assume for a contradiction that $\overrightarrow{T}$ is a tree set of bipartitions of K that distinguishes all the free ultrafilters on K. If $\overrightarrow{T}$ is finite, then there are only finitely many orientations of $\overrightarrow{T}$. But there are infinitely many free ultrafilters on K, so a finite tree set cannot possibly distinguish all of them. Therefore, $\overrightarrow{T}$ must be infinite. Since the empty set does not distinguish any two ultrafilters on K we may assume without loss of generality that $\overrightarrow{T}$ is regular. Then by Lemma 3.5 we know that $\overrightarrow{T}$ contains either an ω-chain or an infinite splitting star.

Suppose first that $\overrightarrow{T}$ contains an ω-chain; that is to say that we find a sequence $(Z_n)_{n<\omega}$ in $\overrightarrow{T}$ with $Z_n \supsetneq Z_{n+1}$ for all n. As $\overrightarrow{T}$ is a tree-set, $K \setminus Z_0$ is non-empty. Put $Z_\omega := \bigcap_{n<\omega} Z_n$. Then Z_ω is nested with every separation in $\overrightarrow{T}$. More precisely, every separation in T has an orientation Z such that either $Z \supseteq Z_n$ for some $n < \omega$ or $Z_\omega \supseteq Z$. We turn the transfinite sequence $(Z_\alpha)_{\alpha \leq \omega}$ into a partition of K, as follows. For every $n < \omega$ set $K_n = Z_n \setminus Z_{n+1}$; and put $K_\omega := (K \setminus Z_0) \cup Z_\omega$. Then $\mathcal{K} := \{K_\alpha \mid \alpha \leq \omega\}$ is an infinite partition of K. Let $\mathcal{U}$ be any free ultrafilter on $\mathcal{K}$, and pick some free ultrafilter U on K with $\mathcal{U} = U.\mathcal{K}$. The free ultrafilter $\mathcal{U}$ contains all cofinite subsets $\{K_m \mid n \leq m < \omega\} \subseteq \mathcal{K}$ with $n < \omega$, and so U contains all $Z_n \setminus Z_\omega$ with $n < \omega$. Recall that every separation in T has an orientation Z such that either $Z \supseteq Z_n$ for some $n < \omega$ or $Z_\omega \supseteq Z$. Hence for every separation $\{Z^*, Z\} \in T$ we have that either $Z \supseteq Z_n$ with $Z_n \setminus Z_\omega \in U$ implies $Z \in U$, or $Z_\omega \supseteq Z$ with $Z_0 \setminus Z_\omega \in U$ implies $Z^* \in U$. Therefore, if $\mathcal{U}'$ is any free ultrafilter on $\mathcal{K}$ other than $\mathcal{U}$, and U' is a free

ultrafilter on K inducing $\mathcal{U}'$, then U' orients every separation in T the same way as U. But then $\vec{T}$ does not distinguish U and U' from each other, a contradiction.

Finally suppose that T contains an infinite splitting star $\sigma = \{ K_i \mid i \in I \}$. If $\mathcal{K} := \{ K_i^* \mid i \in I \}$ is not yet a partition of K, then we add the non-empty interior $\bigcap_{i \in I} K_i$ of σ to $\mathcal{K}$ to turn $\mathcal{K}$ into one. Let $\mathcal{U}$ be any free ultrafilter on $\mathcal{K}$, and pick some free ultrafilter U on K inducing $\mathcal{U}$. The free ultrafilter $\mathcal{U}$ contains all collections $\mathcal{K} - K_i^*$, and hence U contains all K_i. Now every separation in T has an orientation Z with $Z \supseteq K_i$ for some $i \in I$ as σ is splitting, and then $K_i \in U$ implies $Z \in U$. Therefore, if $\mathcal{U}'$ is any free ultrafilter on $\mathcal{K}$ other than $\mathcal{U}$, and U' is a free ultrafilter on K inducing $\mathcal{U}'$, then U' orients every separation in T the same way as $\mathcal{U}$. But then $\vec{T}$ does not distinguish U and U', a contradiction. $\square$

We remark that the proof above even shows the following stronger version of Theorem 3.3: *If K is an infinite set, then for every tree set of bipartitions of K there is a collection of at least $2^{2^{\aleph_0}} = 2^{\mathfrak{c}}$ many free ultrafilters on K all of which induce the same orientation of the tree set.*[3] So if G has precisely one critical vertex set X with $\breve{\mathscr{C}}_X$ countable, then for every tree set of finite order separations of G there is a collection $\mathcal{O}$ of ultrafilter tangles of G such that all ultrafilter tangles in $\mathcal{O}$ induce the same orientation of the tree set and the cardinal $|\mathcal{O}|$ is equal to the total number $2^{\mathfrak{c}}$ of ultrafilter tangles of G.

3.2 The problem case

This subsection is dedicated to examples of graphs whose critical vertex sets give a very bad starting set

$$\big\{ \{X, \breve{\mathscr{C}}_X\} \;\big|\; X \in \mathrm{crit}(G) \big\}.$$

In both cases, all the critical vertex sets interact with each other in a particular way, made precise as follows. Let us say that two critical vertex sets X and Y of G *form a problem case* if X and Y are incomparable as sets and additionally both $C_X(Y) \in \breve{\mathscr{C}}_X$ and $C_Y(X) \in \breve{\mathscr{C}}_Y$ hold.

Example 3.6 Let G be the subtree of the infinitely branching tree $T_{\aleph_0}$ that consists of the first three levels (for an arbitrarily chosen root). Note that the critical vertex sets of any tree are precisely the singletons formed by its infinite degree nodes. Then the collection

$$\big\{ (X, \breve{\mathscr{C}}_X) \;\big|\; X \in \mathrm{crit}(G) \big\}$$

is an infinite star of small separations. As every $\aleph_0$-tangle contains all the small separations (A, V) with A finite (because these can be written as $(A, \mathscr{C}_A)$ and $\mathscr{C}_A \in U(\tau, A)$ for every $\aleph_0$-tangle τ), it follows that every ultrafilter tangle contains this star as a subset, and so no two ultrafilter tangles are distinguished by this star's underlying tree set. Notably, every two distinct critical vertex sets of G form a problem case. $\square$

---

[3] By improving Lemma 3.5 it might be possible to replace $2^{2^{\aleph_0}}$ with $2^{2^{|K|}}$.

Fig. 2 This graph's critical vertex
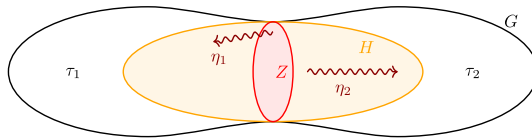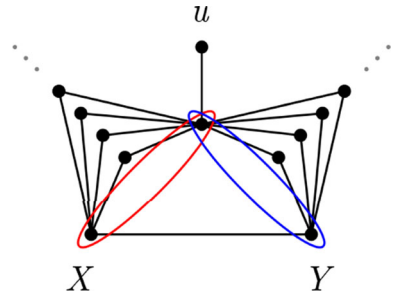sets do not give nested
separations

Fig. 3 The separator Z of a separation efficiently distinguishing two inequivalent tangles τ_1 and τ_2; and a modified torso H with ends η_1 and η_2 representing the two tangles

Example 3.7 If G is the graph shown in Fig. 2, then the collection

$$\{ \, \{X, \breve{\mathscr{C}}_X\} \mid X \in \mathrm{crit}(G) \, \}$$

is not even nested. Indeed, X and Y are the only two critical vertex sets of G. Write V' for $V - u$. Then $\{X, \breve{\mathscr{C}}_X\} = \{X + u, V'\}$ and $\{Y, \breve{\mathscr{C}}_Y\} = \{Y + u, V'\}$. Now these two separations cannot be nested: as X and Y are incomparable as sets, we have neither $(X + u, V') \leq (Y + u, V')$ nor $(V', X + u) \leq (V', Y + u)$. But $(X + u, V') \leq (V', Y + u)$ and $(V', X + u) \leq (Y + u, V')$ are impossible as well since $X + u$ and $Y + u$ are both incomparable with V' as sets. As in the previous example we note that X and Y form a problem case $\square$.

4 The overall proof strategy

Our overall strategy to achieve our main result, Theorem 1, roughly goes as follows. Let G be any infinite connected graph. Recall that every ultrafilter tangle $\tau = (X, U)$ of G naturally comes with a finite-order separation $(X, \breve{\mathscr{C}}_X) \in \tau$. As our first step, we carefully extend and refine the set of these separations into a starting tree set T that already distinguishes all the inequivalent ultrafilter tangles of G, but does not necessarily do so efficiently.

Next, we modify the torsos of T so that every $\aleph_0$-tangle of G is represented in every modified torso by some end of that modified torso. We then show the following assertion (also see Fig. 3): *Let τ_1 and τ_2 be any two inequivalent $\aleph_0$-tangles of G which are not efficiently distinguished by the starting tree set T. For every separator Z efficiently separating the τ_i in G there is a modified torso H of T in which the ends η_i representing the tangles τ_i are efficiently separated by Z.* Now we apply Carmesin's theorem as a black box in all the modified torsos H of T. That is, for every modified torso H of T we obtain a tree set T_H of finite-order separations of H that efficiently distinguishes all the ends of H. Finally, we lift all of Carmesin's tree sets compatibly with each other and with T to obtain a tree set T' of finite-order separations of G that extends T. In the end, every separation in T_H which

efficiently distinguishes two ends η_i in H, with the η_i as in the assertion above, gets lifted to a separation in T' that efficiently distinguishes the τ_i in G.

Phrased differently, we reflect the problem of efficiently distinguishing two inequivalent $\aleph_0$-tangles down to the modified torsos of T. There, the problem reduces to efficiently distinguishing two proxy ends, a problem that has already been solved by Carmesin. Finally, we lift the solutions for the modified torsos of T up to the original graph G to solve the original problem.

5 From principal collections of separators to tree sets

In this section, we show how the separations $\{X, \check{\mathscr{C}}_X\}$ can be slightly modified to give rise to a tree set that comes with quite a list of useful properties. Even though our initial intention is to consider these separations for critical vertex sets X of G, we can prove a much stronger result by more generally considering what we call principal collections of vertex sets:

Definition 5.1   Given a collection $\mathcal{Y}$ of vertex sets of G we say that a vertex set X of G is $\mathcal{Y}$-*principal* if X meets for every $Y \in \mathcal{Y}$ at most one component of $G - Y$. And we say that $\mathcal{Y}$ is *principal* if all its elements are $\mathcal{Y}$-principal.

Notation   If $X \subseteq V(G)$ meets precisely one component of $G - Y$ for some $Y \subseteq V(G)$, then we denote this component by $C_Y(X)$.

Definition 5.2   A set $X \in \mathcal{X}$ is *principal* if it is $\mathcal{X}$-principal.

Example 5.3   An $X \in \mathcal{X}$ is principal, e.g., if it induces a clique $G[X]$ or is included in a critical vertex set of G.

Since principal vertex sets behave like cliques it is possible to alter the graph G so that all principal vertex sets actually induce cliques while the finite-order separations stay the same:

Lemma 5.4   *Suppose that $\mathcal{Y}$ is a collection of principal vertex sets of G and let $G_{\mathcal{Y}}$ be obtained from G by turning each $G[X]$ with $X \in \mathcal{Y}$ into a clique. Then the finite-order separations of G are precisely the finite-order separations of $G_{\mathcal{Y}}$. In particular, $\Theta(G) = \Theta(G_{\mathcal{Y}})$.*

Proof   If $\{A, B\}$ is a finite-order separation of G, then each principal $X \in \mathcal{Y}$ meets at most one component of $G - (A \cap B)$. Therefore, no X adds an $(A \setminus B)$–$(B \setminus A)$ edge in $G_{\mathcal{Y}}$, so $\{A, B\}$ is also a finite-order separation of $G_{\mathcal{Y}}$. The converse holds due to $E(G_{\mathcal{Y}}) \supseteq E(G)$. $\square$

We will use this lemma in Sect. 8 to assume without loss of generality that, for a certain tree set, the torsos coincide with the parts. Our next definition extends 'forming a problem case' from critical vertex sets to arbitrary vertex sets:

Definition 5.5   Two vertex sets X and Y of G with $\{X, Y\}$ principal are said to *form a problem case* if X and Y are incomparable as sets and additionally $C_X(Y) \in \check{\mathscr{C}}_X$ and $C_Y(X) \in \check{\mathscr{C}}_Y$ hold.

The following lemma will keep proofs short:

Lemma 5.6 *If* $\{X, \mathscr{C}\}$ *and* $\{Y, \mathscr{D}\}$ *are separations of G satisfying* $X \cup V[\mathscr{C}] \supseteq Y \cup V[\mathscr{D}]$ *and that each component in* $\mathscr{D}$ *avoids X, then* $(X, \mathscr{C}) \leq (Y, \mathscr{D})$.

Proof It remains to show $V \setminus V[\mathscr{C}] \subseteq V \setminus V[\mathscr{D}]$ which is tantamount to $V[\mathscr{D}] \subseteq V[\mathscr{C}]$, which in turn is evident from the assumptions. $\square$

In the previous section, we have seen that for two distinct critical (in particular principal) vertex sets $X \neq Y$ their separations $\{X, \check{\mathscr{C}}_X\}$ and $\{Y, \check{\mathscr{C}}_Y\}$ need not be nested. This may happen, for example, if X and Y form a problem case. The following two lemmas show that actually this may happen only if X and Y form a problem case.

Lemma 5.7 *Let* $X \subsetneq Y$ *be two vertex sets of G such that* $\{X, Y\}$ *is principal. Then all of the components in* $\check{\mathscr{C}}_Y$ *are properly contained in the component* $C_X(Y)$. *Notably,* $C_X(Y) \in \check{\mathscr{C}}_X$ *if* $\check{\mathscr{C}}_Y$ *is non-empty. Moreover, if we are given subsets* $\mathscr{C} \subseteq \check{\mathscr{C}}_X$ *and* $\mathscr{D} \subseteq \check{\mathscr{C}}_Y$, *then*

$$\vec{S} \leq (X, C_X(Y)) \leq (Y, \check{\mathscr{C}}_Y) \leq (Y, \mathscr{D}) \text{ where } \begin{cases} \vec{S} = (X, \mathscr{C}) & \text{if } C_X(Y) \in \mathscr{C} \\ \vec{S} = (\mathscr{C}, X) & \text{otherwise} \end{cases}$$

so in particular $\{X, \mathscr{C}\}$ *and* $\{Y, \mathscr{D}\}$ *are nested with each other. If additionally the inclusion* $\mathscr{C} \subseteq \check{\mathscr{C}}_X$ *is proper, i.e., if* $(X, \mathscr{C})$ *is not small, then* $(X, \mathscr{C}) \not\leq (\mathscr{D}, Y)$.

Proof Since every component $C \in \check{\mathscr{C}}_Y$ has neighbourhood precisely equal to Y, it follows from $X \subsetneq Y$ that $\bigcup (\check{\mathscr{C}}_Y | X) \subsetneq C_X(Y)$. Hence Lemma 5.6 yields $(X, C_X(Y)) \leq (Y, \check{\mathscr{C}}_Y)$.

It remains to show $(X, \mathscr{C}) \not\leq (\mathscr{D}, Y)$ given that $\mathscr{C} \subsetneq \check{\mathscr{C}}_X$. Assume for a contradiction that the inequality $(X, \mathscr{C}) \leq (\mathscr{D}, Y)$ holds. If $C_X(Y)$ is contained in $\mathscr{C}$, then $\vec{S} = (X, \mathscr{C}) \leq (Y, \mathscr{D})$ is equivalent to $(\mathscr{D}, Y) \leq (\mathscr{C}, X)$ which yields $(X, \mathscr{C}) \leq (\mathscr{D}, Y) \leq (\mathscr{C}, X)$. Hence $(X, \mathscr{C})$ is small which implies $\mathscr{C} = \check{\mathscr{C}}_X$, a contradiction. Otherwise $C_X(Y)$ is not contained in $\mathscr{C}$, and then $\vec{S} = (\mathscr{C}, X) \leq (Y, \mathscr{D})$ is equivalent to $(\mathscr{D}, Y) \leq (X, \mathscr{C})$ which yields $(X, \mathscr{C}) \leq (\mathscr{D}, Y) \leq (X, \mathscr{C})$. In particular, $(X, \mathscr{C}) = (\mathscr{D}, Y)$ implies $X = Y$, contradicting the fact that $X \subsetneq Y$. $\square$

Our next lemma is also illustrated in Fig. 4.

Lemma 5.8 *Let X and Y be two incomparable vertex sets of G such that* $\{X, Y\}$ *is principal. If we are given subsets* $\mathscr{C} \subseteq \check{\mathscr{C}}_X$ *and* $\mathscr{D} \subseteq \check{\mathscr{C}}_Y$ *with* $C_Y(X) \notin \mathscr{D}$, *then*

$$\vec{S} \leq (X, C_X(Y)) \leq (Y, \mathscr{D}) \text{ where } \begin{cases} \vec{S} = (X, \mathscr{C}) & \text{if } C_X(Y) \in \mathscr{C} \\ \vec{S} = (\mathscr{C}, X) & \text{otherwise} \end{cases}$$

so in particular $\{X, \mathscr{C}\}$ *and* $\{Y, \mathscr{D}\}$ *are nested with each other and we have* $(X, \mathscr{C}) \not\leq (\mathscr{D}, Y)$ *(Fig. 4).*
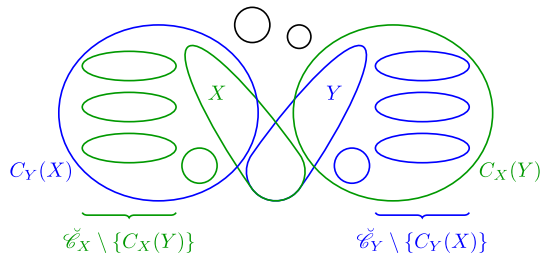
Fig. 4 Two incomparable sets X and Y such that $\{X, Y\}$ is principal. Note that every component of $G - X$ which is neither $C_X(Y)$ nor contained in $C_Y(X)$ has its neighbourhood in $X \cap Y$ and is thus also a component of $G - Y$ (black circles). Also, not every component of $G - Y$ which is contained in $C_X(Y)$ has to be contained in $\breve{\mathscr{C}}_Y$, as is depicted by the blue circle on the right. If $C_Y(X) \notin \mathscr{D} \subseteq \breve{\mathscr{C}}_Y$, then $\mathscr{D}$ is a subset of $\breve{\mathscr{C}}_Y \setminus \{C_Y(X)\}$ and thus $(X, C_X(Y)) \leq (Y, \mathscr{D})$

Proof The assumption $C_Y(X) \notin \mathscr{D}$ ensures that every component in $\mathscr{D}$ avoids X. Let y be any vertex in $Y \setminus X$. As every component in $\mathscr{D}$ avoids X and sends an edge to $y \in Y \setminus X$, we deduce that $(Y \setminus X) \cup \bigcup \mathscr{D} \subseteq C_X(Y)$. Hence Lemma 5.6 yields $(X, C_X(Y)) \leq (Y, \mathscr{D})$. From this, the rest is evident. $\qquad\square$

We are now ready to prove our main technical result, Theorem 5.12. To allow for more flexibility in its applications, we have extracted the following definition and second main technical result from Theorem 5.12:

Definition 5.9 Suppose that $\mathcal{Y}$ is a principal collection of vertex sets of G. A function that assigns to every $X \in \mathcal{Y}$ a subset $\mathscr{K}(X) \subseteq \breve{\mathscr{C}}_X$ is called *admissible* for $\mathcal{Y}$ if for every two $X, Y \in \mathcal{Y}$ that are incomparable as sets we have either $C_X(Y) \notin \mathscr{K}(X)$ or $C_Y(X) \notin \mathscr{K}(Y)$. If additionally $|\breve{\mathscr{C}}_X \setminus \mathscr{K}(X)| \leq 1$ for all $X \in \mathcal{Y}$, then $\mathscr{K}$ is *strongly* admissible for $\mathcal{Y}$.

Theorem 5.10 *For every principal collection of vertex sets of a connected graph there is a strongly admissable function.*

Proof Let $\mathcal{Y}$ be a principal collection of vertex sets of a connected graph G. We write $\mathscr{P}$ for the collection of those principal vertex sets in $\mathcal{Y}$ that form a problem case with some other principal vertex set in $\mathcal{Y}$. Let us fix any well-ordering of $\mathscr{P}$ and view $\mathscr{P}$ as well-ordered set from now on.

For each $X \in \mathscr{P}$ we put $K(X) := C_X(Y)$ for the first $Y \in \mathscr{P}$ which forms a problem case with X. Let us put $\mathscr{K}(X) := \breve{\mathscr{C}}_X \setminus \{K(X)\}$ for every $X \in \mathscr{P}$, and $\mathscr{K}(X) := \breve{\mathscr{C}}_X$ for all other vertex sets $X \in \mathcal{Y}$. We claim that $\mathscr{K}$ is strongly admissible for $\mathcal{Y}$.

For this, let $X \neq Y$ be any two distinct vertex sets in $\mathcal{Y}$ that form a problem case. We show that at least one of $K(X) = C_X(Y)$ and $K(Y) = C_Y(X)$ holds. Let $Z \in \mathscr{P}$ be the first vertex set that forms a problem case with one of X and Y. Without loss of generality we may assume that Z forms a problem case with X, so we have $K(X) = C_X(Z)$ by the minimal choice of Z. Since we are done if Y and Z meet the same component of $G - X$, we may assume that $C_X(Y) \neq C_X(Z)$. This means that the three sets X, Y, Z are pairwise incomparable. Our plan is to show that Y forms a problem case with Z, and that this gives $K(Y) = C_Y(Z) = C_Y(X)$ as desired.

We already know that Y and Z are incomparable. Next, let us verify that $C_Y(Z) \in \breve{\mathscr{C}}_Y$. For this, pick any vertex $x \in X \setminus Y$. As X and Z form a problem case we have $C_X(Z) \in \breve{\mathscr{C}}_X$, so the vertex x sends some edge e to the component $C_X(Z)$. Now x is not in Y and the component

$C_X(Z)$ avoids Y as Y and Z live in distinct components of $G - X$ by assumption, so $C_X(Z) + e$ is a connected subgraph of $G - Y$ that meets both X and Z, yielding $C_Y(Z) = C_Y(X)$. Since Y and X form a problem case, giving $C_Y(X) \in \breve{\mathscr{C}}_Y$, we get $C_Y(Z) \in \breve{\mathscr{C}}_Y$ as required. By symmetry we have $C_Z(Y) \in \breve{\mathscr{C}}_Z$, so Y and Z form a problem case as desired and $K(Y) = C_Y(Z)$ follows from the minimal choice of Z. To see that $K(Y) = C_Y(X)$ holds, recall that we proved $C_Y(Z) = C_Y(X)$ three sentences earlier. $\qquad\square$

Remark 5.11 Admissible functions are monotone in the following sense. If $\mathscr{K}$ is an admissible function for a principal collection $\mathcal{Y}$ of vertex sets of a graph, then every function $\mathscr{K}'$ that assigns to each $X \in \mathcal{Y}$ a subset $\mathscr{K}'(X) \subseteq \mathscr{K}(X)$ is also admissible. For instance, if $\mathcal{Y}$ is a collection of critical vertex sets, then Theorem 5.10 yields a strongly admissible function $\mathscr{K}$ for $\mathcal{Y}$, and by discarding arbitrarily chosen components from some component collections $\mathscr{K}(X)$ we may see to it that $|\breve{\mathscr{C}}_X \setminus \mathscr{K}(X)| = 1$ for all $X \in \mathcal{Y}$. We will do this in later proofs to ensure that the separations of the form $\{X, \mathscr{K}(X)\}$ have no small orientations, which we will need to obtain tree sets from them.

Finally, we go for Theorem 5.12, which considers tree sets of the following form:

Notation Given a principal collection $\mathcal{Y}$ of vertex sets of G and an admissible function $\mathscr{K}$ for $\mathcal{Y}$ we write

$$T(\mathcal{Y}, \mathscr{K}) := \big\{ \{X, \mathscr{K}(X)\}, \{X, K\} \mid X \in \mathcal{Y} \text{ and } K \in \mathscr{K}(X) \big\}.$$

For every vertex set $X \in \mathcal{Y}$ we write $\sigma_X^{\mathscr{K}}$ for the star that consists of the separation $(X, \mathscr{K}(X))$ and all the separations (K, X) with $K \in \mathscr{K}(X)$. Notably, each star $\sigma_X^{\mathscr{K}}$ has interior X.

For an illustration of the following theorem, see Fig. 5.

Theorem 5.12 *Let G be any connected graph, let $\mathcal{Y}$ be a principal collection of vertex sets of G and let $\mathscr{K}$ be an admissible function for $\mathcal{Y}$. Abbreviate $T(\mathcal{Y}, \mathscr{K}) = T$ and $\sigma_X^{\mathscr{K}} = \sigma_X$. Then the following assertions hold:*

Fig. 5 A principal set $\mathcal{Y} = \{W, X, Y, Z\}$ of pairwise disjoint sets and the separations of the form $(\mathscr{K}(X'), X')$ for $X' \in \mathcal{Y}$ where $\mathscr{K}$ is some admissible function for $\mathcal{Y}$. Note that in accordance with part (i) of Theorem 5.12 the depicted separations form a partial consistent orientation
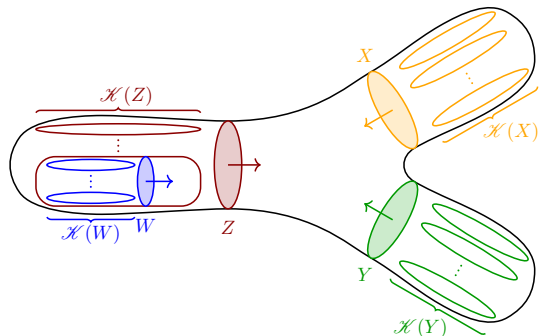
(i) *For every distinct two $X, Y \in \mathcal{Y}$, after possibly swapping X and Y, either*

$$(\mathcal{K}(X), X) \leq (Y, \mathcal{K}(Y)) \text{ or } (X, \mathcal{K}(X)) \leq (X, C_X(Y)) \leq (Y, \mathcal{K}(Y)).$$

The collection of all separations $(\mathcal{K}(X), X)$ with $\mathcal{K}(X) \subsetneq \mathscr{C}_X$ forms a consistent partial orientation of T.

(ii) *The collection T of separations is nested. It is a regular tree set if $\emptyset \subsetneq \mathcal{K}(X) \subsetneq \mathscr{C}_X$ holds for all $X \in \mathcal{Y}$.*

(iii) *Every star σ_X with $X \in \mathcal{Y}$ is a splitting star of $\overrightarrow{T}$.*

Moreover, if all the vertex sets in $\mathcal{Y}$ are finite, then we may add:

(iv) *If τ is an ultrafilter tangle of G with $X_\tau \in \mathcal{Y}$ and $\mathcal{K}(X_\tau) \in U(\tau, X_\tau)$, then τ induces via $\tau \mapsto \tau \cap \overrightarrow{T}$ on T the consistent orientation which is given by the infinite splitting star σ_{X_τ} in that $\tau \cap \overrightarrow{T} = \lceil \sigma_{X_\tau} \rceil$.*

(v) *If $\mathrm{crit}(G) \subseteq \mathcal{Y}$ and $\check{\mathscr{C}}_X \setminus \mathcal{K}(X)$ is finite for all $X \in \mathrm{crit}(G)$, then T distinguishes every two inequivalent ultrafilter tangles τ_1 and τ_2 of G via separations in $\sigma_{X_{\tau_1}}$ and $\sigma_{X_{\tau_2}}$, and it distinguishes every end from every ultrafilter tangle τ via a separation in σ_{X_τ}.*

Proof

(i) We start with the inequalities. If X and Y are comparable with $X \subsetneq Y$, say, then we are done by Lemma 5.7. Otherwise X and Y are incomparable, and then we are done by Lemma 5.8 since $\mathcal{K}$ is admissable. The two lemmas also prevent $(X, \mathcal{K}(X)) \leq (\mathcal{K}(Y), Y)$ for all distinct two elements $X, Y \in \mathcal{Y}$ with $\mathcal{K}(X) \subsetneq \mathscr{C}_X$ and $\mathcal{K}(Y) \subsetneq \mathscr{C}_Y$, ensuring that the partial orientation of T formed by the separations $(\mathcal{K}(X), X)$ with $\mathcal{K}(X) \subsetneq \mathscr{C}_X$ is consistent.

(ii) That T is nested follows from (i). For the 'moreover' part note that requiring $\emptyset \subsetneq \mathcal{K}(X) \subsetneq \mathscr{C}_X$ ensures that $\{X, \mathcal{K}(X)\}$ has no small orientation.

(iii) It suffices to show that every separation in T with separator $Y \neq X$ has an orientation that lies below some element of σ_X. So consider any $Y \in \mathcal{Y}$ other than X. Since σ_Y is a star, it suffices to show that some separation in $(\sigma_Y)^*$ lies below some element of σ_X. By (i) it suffices to consider the following cases. If $(\mathcal{K}(Y), Y) \leq (X, \mathcal{K}(X))$ we are done. Otherwise either

$$(X, \mathcal{K}(X)) \leq (X, C_X(Y)) \leq (Y, \mathcal{K}(Y)) \text{ or } (Y, \mathcal{K}(Y)) \leq (Y, C_Y(X)) \leq (X, \mathcal{K}(X)).$$

In the first case we are fine since $(\mathcal{K}(Y), Y) \leq (C_X(Y), X) \in \sigma_X$. And in the second case we are done by the second inequality.

(iv) Let τ be any ultrafilter tangle of G with $X_\tau \in \mathcal{Y}$ and write $X = X_\tau$. First, we show that σ_X is included in $O := \tau \cap \overrightarrow{T}$. The assumption $\mathcal{K}(X) \in U(\tau, X)$ means $(X, \mathcal{K}(X)) \in O$. Moreover, we have $(K, X) \in \tau$ for every $K \in \mathcal{K}(X_\tau)$ as $U(\tau, X)$ is a free ultrafilter. Thus $\sigma_X \subseteq O$, and so $\lceil \sigma_X \rceil \subseteq O$ by consistency. Conversely, $O \subseteq \lceil \sigma_X \rceil$ since $\sigma_X \subseteq O$ is a splitting star of $\overrightarrow{T}$ by (iii).

(v) If τ_1 and τ_2 are two ultrafilter tangles of G with $X_{\tau_1} \neq X_{\tau_2}$, then the induced orientations $\tau_i \cap \overrightarrow{T}$ come from distinct splitting stars $\sigma_{X_{\tau_i}}$ of $\overrightarrow{T}$ by (iii). Now if ω is

an end of G and τ is an ultrafilter tangle, then ω avoids the star σ_{X_τ} since it has finite interior (cf. [4, Corollary 1.7]) while τ contains it by (iii).

$\square$

We close this section by showing that in general it is not possible to find an admissable function $\mathscr{H}$ for which $\vec{T}(\mathrm{crit}(G), \mathscr{H})$ is a tree set that is even isomorphic to the edge tree set of a tree.

Example 5.13 If G is the graph shown in Fig. 6, then there is no function assigning to each critical vertex set X of G a cofinite subset $\mathscr{H}(X) \subseteq \breve{\mathscr{C}}_X$ such that

$$N := \big\{ \{X, \mathscr{H}(X)\} \mid X \in \mathrm{crit}(G) \big\}$$

gives rise to a tree set $\vec{N}$ that is isomorphic to the edge tree set of a tree (so in particular it cannot be induced by an $S_{\aleph_0}$-tree or tree-decomposition of G). First, however, we describe G more precisely. The graph G is obtained from the $\aleph_0$-regular tree $T = T_{\aleph_0}$ by fixing any root r and then proceeding as follows. For every ray $R \subseteq T$ starting at the root r we add a new copy of $K_{2,\aleph_0}$ with 2-class $\{x_R, y_R\}$, say, and join x_R to every vertex of the ray R. Readers familiar with the 'binary tree with tops' will note that G extends a '$T_{\aleph_0}$ with tops'.

Let us check that there really is no suitable function $X \mapsto \mathscr{H}(X)$ as claimed. Assume for a contradiction that there is. Then $\vec{N}$ is a tree set that, by Theorem 2.2, has no $(\omega + 1)$-chains. Hence to yield a contradiction, it suffices to find an $(\omega + 1)$-chain. If t is a node of $T \subseteq G$, then its down-closure $\lceil t \rceil$ in T is a critical vertex set of G, and the components in $\breve{\mathscr{C}}_{\lceil t \rceil}$ are of the following form. If t' is an upward neighbour of t in T, then the vertex set of the component of $G - \lceil t \rceil$ containing t' is given by the union of $\lfloor t' \rfloor \subseteq T$ with all the copies of $K_{2,\aleph_0}$ whose corresponding ray has a tail in $\lfloor t' \rfloor$. This gives a bijection between the upward neighbours of t in T and the components in $\breve{\mathscr{C}}_{\lceil t \rceil}$. Next, we claim that there is a ray $R^* = t_0 t_1 t_2 \ldots \subseteq T$ starting at the root r such that for all $n > 0$ the node t_n corresponds to a component in $\mathscr{H}(\lceil t_{n-1} \rceil)$ for its predecessor t_{n-1}. Indeed, since $\mathscr{H}(\lceil t \rceil) \subseteq \breve{\mathscr{C}}_{\lceil t \rceil}$ is infinite for all $t \in T$, such a ray can be constructed inductively. But then we get a strictly ascending sequence

$$(\lceil t_0 \rceil, \mathscr{H}(\lceil t_0 \rceil)) < (\lceil t_1 \rceil, \mathscr{H}(\lceil t_1 \rceil)) < (\lceil t_2 \rceil, \mathscr{H}(\lceil t_2 \rceil)) < \cdots$$

i.e. we get an ω-chain in $\vec{N}$. And this ω-chain extends to an $(\omega + 1)$-chain as the separation

Fig. 6 A $T_{\aleph_0}$ (black) with $2^{\aleph_0}$ many copies of $K_{2,\aleph_0}$ as 'tops' (visualised in red for the rightmost ray)
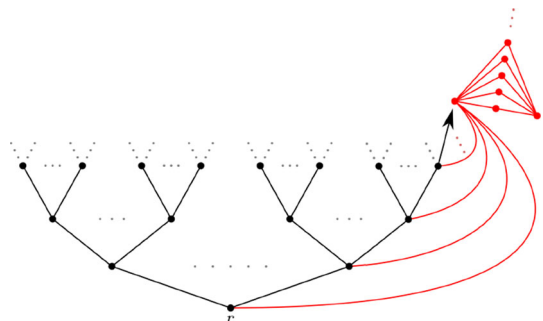
$(Z, \mathscr{K}(Z))$ with $Z = \{x_{R^*}, y_{R^*}\}$ that comes from the $K_{2,\aleph_0}$ for R^* is greater than all separations $(\lceil t_n \rceil, \mathscr{K}(\lceil t_n \rceil))$.

6 Applications

This section is dedicated to the applications of our work mentioned in the introduction. All of the four applications are, in fact, applications of Theorems 5.10 and 5.12. Elbracht, Kneip and Teegen [9] use the following corollary of our two theorems:

Corollary 6.1 *Suppose that $\mathcal{Y}$ is a principal[4] collection of vertex sets of G. Then there is a function $\mathscr{K}$ assigning to each $X \in \mathcal{Y}$ a subset $\mathscr{K}(X) \subseteq \check{\mathscr{C}}_X$ (the set $\check{\mathscr{C}}_X$ consists of the components of $G - X$ whose neighbourhoods are precisely equal to X) that misses at most one component from $\check{\mathscr{C}}_X$, such that the collection*

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

is nested.

Bürger and the second author [10] use Theorems 5.10 and 5.12 directly. In the remainder of this section, we present applications three and four: a structural connectivity result for infinite graphs, and the collectionwise normality of ultrafilter tangle spaces.

6.1 A structural connectivity result for infinite graphs

We have already explained this application in detail in our introduction, now we prove it:

Theorem 2 *Every connected graph G has a tree set whose separators are precisely the critical vertex sets of G and all whose torsos are tough.*

Proof By Theorems 5.10 and 5.12 it suffices to show that for $\mathcal{Y} := \mathrm{crit}(G)$ and a strongly admissable function $\mathscr{K}$ the torsos of the tree set $T(\mathcal{Y}, \mathscr{K})$ are tough. For this, let O be any consistent orientation of $T(\mathcal{Y}, \mathscr{K})$, let Π be its part and H its torso. In order to show that H is tough, let Ξ be a finite subset of $V(H)$. Let $\mathscr{C} \subseteq \mathscr{C}_\Xi$ consist of those components of $G - \Xi$ that meet Π. Then $\mathscr{C}$ must be finite: otherwise Ξ contains a critical vertex set Ξ' of G with $\mathscr{C}' := \check{\mathscr{C}}_{\Xi'} \cap \mathscr{C}$ infinite; and then $(\Xi', C) \in O$ for all $C \in \mathscr{C}' \cap \mathscr{K}(\Xi')$ as these C meet Π, contradicting the consistency of O. Thus $G - \Xi$ has only finitely many components meeting Π. By Corollary 2.6 each of these components induces a component of $H - \Xi$, and so $H - \Xi$ has only finitely many components. $\square$

6.2 Collectionwise normality of the ultrafilter tangle space

For this subsection, we recall the following definitions from general topology (cf. [24]):

---

[4] cf. Definition 5.1

Definition 6.2 (Normality and collectionwise normality) Let X be a topological space in which all singletons are closed.

The space X is said to be *normal* if for every two disjoint closed subsets A_1 and A_2 of X there are disjoint open subsets O_1 and O_2 of X with $A_1 \subseteq O_1$ and $A_2 \subseteq O_2$.

A collection $\{ A_i \mid i \in I \}$ of subsets $A_i \subseteq X$ is *discrete* if every point $x \in X$ has an open neighbourhood that meets at most one of the A_i.

The space X is said to be *collectionwise normal* if for every discrete collection $\{ A_i \mid i \in I \}$ of pairwise disjoint closed subsets $A_i \subseteq X$ there exists a collection $\{ O_i \mid i \in I \}$ of pairwise disjoint open subsets $O_i \subseteq X$ with $A_i \subseteq O_i$ for all $i \in I$.

The following implications are true for every topological space (the first implication is [24, Theorems 5.1.1 and 5.1.18] whereas the second is clear):

$$\text{compact Hausdorff} \;\Rightarrow\; \text{collectionwise normal} \;\Rightarrow\; \text{normal.}$$

In general, the end space $\Omega(G)$ of a graph is not compact, e.g., if G is a union of infinitely many rays sharing precisely their first vertex. Polat [12] and Sprüssel [13] independently showed that the end space of every graph G is normal, and Polat even showed that the end space is collectionwise normal (this is Lemma 4.14 in [12], see [25] for a modern proof):

Theorem *Every graph G has a collectionwise normal end space $\Omega(G)$.*

The $\aleph_0$-tangle space $\Theta(G)$, with the subspace topology imposed by Diestel's tangle compactification, is compact Hausdorff and contains the end space as a subspace. Since every compact Hausdorff space is collectionwise normal, the $\aleph_0$-tangle space is collectionwise normal as well:

Theorem *Every graph G has a collectionwise normal $\aleph_0$-tangle space $\Theta(G)$.*

This result, however, does not imply that the end space is collectionwise normal, for usually the end space is not closed in the $\aleph_0$-tangle space.

As the $\aleph_0$-tangle space is the disjoint union $\Theta(G) = \Omega(G) \sqcup \Upsilon(G)$ of the end space $\Omega(G)$ and the ultrafilter tangle space $\Upsilon(G)$, the question arises whether the ultrafilter tangle space is collectionwise normal as well. Like the end space, the ultrafilter tangle space usually is not closed in the $\aleph_0$-tangle space, so the ultrafilter tangle space does not obviously inherit the collectionwise normality from the $\aleph_0$-tangle space.

In this section, we show that the ultrafilter tangle space is collectionwise normal, Theorem 3. We remark that our proof also shows that the critical vertex set space (with the subspace topology from the compactification $|G|_\Gamma = G \sqcup \mathrm{crit}(G) \sqcup \Omega(G)$ introduced in [6]) is collectionwise normal as well.

Theorem 3 *Let G be any connected graph. Then for every discrete collection $\{ A_i \mid i \in I \}$ of pairwise disjoint closed subsets $A_i \subseteq \Upsilon(G)$ there exists a collection $\{ O_i \mid i \in I \}$ of pairwise disjoint open subsets $O_i \subseteq |G|_\Theta$ such that $A_i \subseteq O_i$ for all $i \in I$. In particular, the ultrafilter tangle space of G is collectionwise normal.*

Our proof of Theorem 3 will employ the following short lemma:

Lemma 6.3 *For every two finite-order separations* $(X, \mathscr{C}) \leq (Y, \mathscr{D})$ *of* G *we have* $\mathcal{O}_{|G|_{\Theta}}(X, \mathscr{C}) \supseteq \mathcal{O}_{|G|_{\Theta}}(Y, \mathscr{D})$.

Proof Clearly, $G \cap \mathcal{O}_{|G|_{\Theta}}(X, \mathscr{C}) \supseteq G \cap \mathcal{O}_{|G|_{\Theta}}(Y, \mathscr{D})$. And from the consistency of $\aleph_0$-tangles we deduce $\Theta \cap \mathcal{O}_{|G|_{\Theta}}(X, \mathscr{C}) \supseteq \Theta \cap \mathcal{O}_{|G|_{\Theta}}(Y, \mathscr{D})$. $\qquad\square$

Proof of Theorem 3 For this, let $\{A_i \mid i \in I\}$ by any discrete collection of closed subsets $A_i \subseteq \Upsilon(G)$. We are going to find a suitable collection $\{O_i \mid i \in I\}$. To get started, we view the $\aleph_0$-tangle space as inverse limit $\Theta = \varprojlim(\beta(\mathscr{C}_X) \mid X \in \mathcal{X})$ where each $\mathscr{C}_X$ is endowed with the discrete topology. Since Θ is compact and all $\beta(\mathscr{C}_X)$ are Hausdorff, it follows from general topology that all of the (continuous) projections $\mathrm{pr}_Y : \Theta = \varprojlim\beta(\mathscr{C}_X) \to \beta(\mathscr{C}_Y)$ are closed. Now consider any critical vertex set X of G. The Stone-Čech remainder $(\mathscr{C}_X)^* = \beta(\check{\mathscr{C}}_X) \setminus \check{\mathscr{C}}_X$ is a closed subspace of $\beta(\check{\mathscr{C}}_X) = \mathrm{cl}_{\beta(\check{\mathscr{C}}_X)}(\check{\mathscr{C}}_X) \subseteq \beta(\check{\mathscr{C}}_X)$. (This follows from general topology, but it can also be seen more directly by considering the standard basis for the Stone-Čech compactification of discrete spaces.) Since pr_X and $(\check{\mathscr{C}}_X)^*$ are closed, for every $i \in I$ the set

$$A_{i,X} := \mathrm{pr}_X(A_i) \cap (\check{\mathscr{C}}_X)^* = \mathrm{pr}_X(\overline{A_i}^{\Theta}) \cap (\check{\mathscr{C}}_X)^*$$

is closed in $\beta(\check{\mathscr{C}}_X)$. Note that the sets $A_{i,X}$ may be empty. We claim that $\{A_{i,X} \mid i \in I\}$ is a discrete collection of pairwise disjoint closed subsets of $\beta(\check{\mathscr{C}}_X)$. For this, consider any $U \in (\check{\mathscr{C}}_X)^*$. The preimage $\mathrm{pr}_X^{-1}(U)$ is a singleton that consists precisely of the ultrafilter tangle τ of which (X, U) is the blueprint. Since the sets A_i are pairwise disjoint, this means that the sets $A_{i,X}$ must be pairwise disjoint as well. If τ is contained in a set A_i, then τ has an open neighbourhood in $\Upsilon(G)$ which avoids all sets A_j with $j \neq i$, because $\{A_i \mid i \in I\}$ is discrete. In particular, τ has an open neighbourhood $\mathcal{O}_{|G|_{\Theta}}(Y, \mathscr{D})$ in $|G|_{\Theta}$ which avoids all A_j with $j \neq i$. By the 'in particular' part of Lemma 3.1 and by Lemma 6.3, we may assume that $Y = X$ and $\mathscr{D} \in U$. Hence the set $\{U' \in \beta(\check{\mathscr{C}}_X) \mid \mathscr{D} \in U'\}$ induced by $\mathcal{O}_{|G|_{\Theta}}(X, \mathscr{D})$ is an open neighbourhood of U in $\beta(\check{\mathscr{C}}_X)$ which avoids all $A_{j,X}$ with $j \neq i$. Therefore, $\{A_{i,X} \mid i \in I\}$ is a discrete collection of pairwise disjoint closed subsets of $\beta(\check{\mathscr{C}}_X)$. Now the Stone-Čech compactification $\beta(\check{\mathscr{C}}_X)$ is collectionwise normal since it is compact Hausdorff, and so we find a collection $\{O_{i,X} \mid i \in I\}$ of pairwise disjoint open subsets $O_{i,X} \subseteq \beta(\check{\mathscr{C}}_X)$ satisfying the inclusion $A_{i,X} \subseteq O_{i,X}$ for all $i \in I$.

Next, we use Theorem 5.10 and Remark 5.11 to find a strongly admissable function $\mathscr{K}$ for the principal collection $\mathrm{crit}(G)$ with $|\check{\mathscr{C}}_X \setminus \mathscr{K}(X)| = 1$ for all $X \in \mathrm{crit}(G)$. For every index $i \in I$ and every ultrafilter tangle $\tau \in A_i$ we choose a component collection $\mathscr{C}(\tau) \in U(\tau, X_\tau)$ such that

- $\mathscr{C}(\tau) \subseteq \mathscr{K}(X_\tau)$;
- $\mathscr{C}(\tau) \subseteq O_{i,X_\tau}$;
- $O_{i,\tau} := \mathcal{O}_{|G|_{\Theta}}(X_\tau, \mathscr{C}(\tau))$ avoids all A_j with $j \neq i$.

We find $\mathscr{C}(\tau)$ as follows. First, we recall that $\mathscr{K}(X_\tau)$ is contained in the free ultrafilter $U(\tau, X_\tau)$. Second, we note that $O_{i,X_\tau} \cap \check{\mathscr{C}}_{X_\tau}$ is contained in $U(\tau, X_\tau)$ as well, for O_{i,X_τ} is an open neighbourhood of $U = \mathrm{pr}_{X_\tau}(\tau) \in A_{i,X_\tau}$ in $\beta(\check{\mathscr{C}}_{X_\tau})$ and U is contained in $U(\tau, X_\tau)$ as a subset. Therefore, if we find a component collection $\mathscr{C} \subseteq \check{\mathscr{C}}_{X_\tau}$ with $\mathscr{C} \in U(\tau, X_\tau)$ such that

$\mathscr{O}_{|G|_\Theta}(X_\tau, \mathscr{C})$ avoids all A_j with $j \neq i$, then $\mathscr{C}(\tau) := \mathscr{K}(X_\tau) \cap O_{i,X_\tau} \cap \mathscr{C}$ will satisfy all three requirements (for the third requirement we apply Lemma 6.3 to $(X_\tau, \mathscr{C}) \leq (X_\tau, \mathscr{C}(\tau))$). To find a suitable component collection $\mathscr{C}$, we proceed as follows. The union of all sets A_j with $j \in I$ and $j \neq i$ is closed in $\Upsilon(G)$ since $\{ A_i \mid i \in I \}$ is a discrete collection of closed sets. Hence there exists an open neighbourhood $\mathscr{O}_{|G|_\Theta}(Y, \mathscr{D})$ of τ in $|G|_\Theta$ which avoids this union. Applying Lemma 3.1 to $(Y, \mathscr{D}) \in \tau$ then yields a component collection $\mathscr{C} \subseteq \check{\mathscr{C}}_{X_\tau}$ satisfying $(Y, \mathscr{D}) \leq (X_\tau, \mathscr{C}) \in \tau$. In particular, $\mathscr{O}_{|G|_\Theta}(X_\tau, \mathscr{C}) \subseteq \mathscr{O}_{|G|_\Theta}(Y, \mathscr{D})$ (Lemma 6.3 again) avoids all A_j with $j \neq i$.

Letting $O_i := \bigcup \{ O_{i,\tau} \mid \tau \in A_i \}$ for every $i \in I$, we claim that the collection $\{ O_i \mid i \in I \}$ is as desired. For this, it suffices to show that for all indices $i \neq j$ and ultrafilter tangles $\tau \in A_i$ and $\tau' \in A_j$ the open neighbourhoods $O_{i,\tau}$ and $O_{j,\tau'}$ are disjoint. By Theorem 5.12 (i) and by symmetry, only the following three cases can possibly occur.

In the first case we have $X_\tau = X_{\tau'}$ and write $X = X_\tau$. Then $O_{i,X}$ and $O_{j,X}$ are disjoint, ensuring that $\mathscr{C}(\tau)$ and $\mathscr{C}(\tau')$ are disjoint. (If we had not involved the open sets $O_{i,X}$ and $O_{j,X}$, then the component collections $\mathscr{C}(\tau)$ and $\mathscr{C}(\tau')$ might possibly have a non-empty finite intersection.) In particular, $O_{i,\tau}$ and $O_{j,\tau'}$ are disjoint as well.

In the second case we have $X_\tau \neq X_{\tau'}$ and $(\mathscr{K}(X_\tau), X_\tau) \leq (X_{\tau'}, \mathscr{K}(X_{\tau'}))$, which implies that $O_{i,\tau}$ and $O_{j,\tau'}$ are disjoint.

In the third case we have $X_\tau \neq X_{\tau'}$ and $(X_\tau, \mathscr{K}(X_\tau)) \leq (X_\tau, C) \leq (X_{\tau'}, \mathscr{K}(X_{\tau'}))$ where C is the component $C_{X_\tau}(X_{\tau'})$. Since $O_{i,\tau}$ avoids $A_j \ni \tau'$ we deduce that the component C is not contained in $\mathscr{C}(\tau)$. Hence $(\mathscr{C}(\tau), X_\tau) \leq (X_{\tau'}, \mathscr{C}(\tau'))$ which implies that $O_{i,\tau}$ and $O_{j,\tau'}$ are disjoint. $\qquad\square$

7 Consistent orientation and lifting from torsos

For this section, fix a graph G, a regular tree set N of finite-order separations of G, and a consistent orientation O of N. Also define $\Pi = \bigcap_{(C,D) \in O} D$.

This section deals with the problem of translating separations of $\mathrm{torso}(G, O)$ to separations of G, as described in Sect. 4. More precisely, given a separation (A, B) of $\mathrm{torso}(G, O)$, we want to find an *extension* of it in G, a separation (U, W) of G towards which all elements of O point such that $U \cap W \subseteq \Pi$ and $(U \cap \Pi, W \cap \Pi) = (A, B)$. Note that every extension (U, W) of (A, B) satisfies $U \cap W = A \cap B$. In general, extensions are not unique. However, the information contained in O already puts strong restrictions on the structure of extensions.

On the one hand, if x and y are vertices of G and (C, D) is a separation in O with $\{x, y\} \subseteq C$ then every extension (U, W) of a separation of $\mathrm{torso}(G, O)$ has to satisfy $(C, D) \leq (U, W)$ or $(C, D) \leq (W, U)$ and thus $\{x, y\} \subseteq U$ or $\{x, y\} \subseteq W$. So here we have a relation on $\bigcup_{(C,D) \in O} C$ and related vertices cannot be separated by extensions of separations of $\mathrm{torso}(G, O)$.

On the other hand, if (C, D) and (C', D') are separations in O such that O also contains some (C'', D'') with $(C, D) \leq (C'', D'')$ and $(C', D') \leq (C'', D'')$, then (C, D) and (C', D') cannot lie on different sides of (U, W) because (C'', D'') points towards every extension (U, W) of (A, B). So here we have a relation on O and no extension of a separation of $\mathrm{torso}(G, O)$ can separate two related separations in O.

It turns out that the two relations describe two points of view on the same idea: In this paper we define $\sim$ as a relation on the set of separations of O, as that fits better in our

framework of tree sets. But it is possible just as well to work with the relation on vertices, as is done e.g. in [3], and several lemmas in this section are inspired by similar lemmas in that paper. Indeed, we will associate with every equivalence class γ of $\sim$ a set of vertices A_γ, thereby associating an equivalence class of $\sim$ of separations with an equivalence class of vertices, and we will work with both γ and A_γ.

Lemma 7.1 *Define a relation $\sim$ on O where $(C,D) \sim (C',D')$ if and only if there is a separation in O above both (C, D) and (C',D'). Then $\sim$ is an equivalence relation.*

Proof By definition the relation is reflexive and symmetric. In order to show transitivity, assume that (C, D) and (C',D') are related, as witnessed by $(U, W) \in O$, and that (C',D') and (C'',D'') are related, as witnessed by $(U',W') \in O$. As O is a consistent orientation, we have $(U, W) \leq (U',W')$ or $(W',U') \leq (W,U)$ or $(U, W) \leq (W',U')$. But $(U, W) \leq (W',U')$ implies $(C',D') \leq (U, W) \leq (W',U') \leq (D',C')$ and thus that $(C',D') \leq (D',C')$ which contradicts the fact that N is regular. So either $(U, W) \leq (U',W')$ or $(U',W') \leq (U, W)$ and in both cases the bigger one of these separations shows that (C, D) and (C'',D'') are related. □

Definition 7.2 An equivalence class of the relation from Lemma 7.1 is a *corridor* of O. For a corridor γ let A_γ be the union of all sets C where $(C,D) \in \gamma$.

Remark 7.3 Let γ be a corridor and (A, B) the supremum of all elements of γ. Then $A = A_\gamma$ and $A \cap B = A_\gamma \cap \Pi$.

Remark 7.4 Lemma 7.1 also holds in abstract separation systems with the same proof. In particular, corridors are well-defined for abstract separation systems.

Lemma 7.5 *If (C, D) and (C',D') are elements of O and $C \setminus D'$ is non-empty then (C, D) and (C',D') are comparable.*

Proof Because O is consistent and nested, any two separations in O either point towards each other or are comparable. Let w be a vertex contained in $C \setminus D'$. Then w witnesses that $(C,D) \not\leq (D',C')$, hence (C, D) and (C',D') do not point towards each other. □

Lemma 7.6 *Let γ be a corridor of O and U a finite subset of A_γ. Then there is a separation (C, D) in γ such that C contains U and $C \setminus D$ contains $U \setminus \Pi$.*

Proof First we consider the special case that U contains only one vertex $v \notin \Pi$. As v is a vertex of A_γ there is a separation (C, D) in γ such that $v \in C$. Furthermore because v is not contained in Π there is a separation (C',D') in O such that v is contained in $C' \setminus D'$. By Lemma 7.5 the separations (C, D) and (C',D') are comparable and thus contained in the same corridor, so (C',D') is contained in γ.

Now consider an arbitrary finite subset U of A_γ. For every vertex v of U there is a separation (C_v,D_v) in γ such that $v \in C_v$. We just showed that if v is not contained in Π then (C_v,D_v) can be chosen such that D_v does not contain v. As γ is a corridor and U is finite, there is a separation (C, D) in γ which is bigger than or equal to all separations (C_v,D_v). In particular $v \in C_v \subseteq C$ for all $v \in U$ and $v \in C_v \setminus D_v \subseteq C \setminus D$ for all $v \in U \setminus \Pi$. □

Lemma 7.7 *The sets $A_\gamma \setminus \Pi$ partition $V(G) \setminus \Pi$.*

Proof By definition of Π every vertex $v \in V(G) \setminus \Pi$ is contained in A_γ for some corridor γ, so the sets $A_\gamma \setminus \Pi$ cover $V(G) \setminus \Pi$. To prove their disjointness, assume that some vertex is contained in A_γ and $A_{\gamma'}$ for two corridors γ and γ'. By Lemma 7.6 there are separations (C, D) in γ and (C', D') in γ' respectively such that both $C \setminus D$ and $C' \setminus D'$ contain v. Thus by Lemma 7.5 the separations (C, D) and (C', D') are contained in the same corridor and hence $\gamma = \gamma'$. □

Corollary 7.8 *A separation (C, D) of N is contained in a given corridor γ if and only if $C \setminus D \subseteq A_\gamma$.*

Lemma 7.9 *Let U be a connected set of vertices avoiding Π. Then there is a corridor γ with $U \subseteq A_\gamma$.*

Proof By Lemma 7.7 it is sufficient to show the statement for U with exactly two elements. Let v and w be two neighbours not in Π, and let (C, D) be a separation in O such that $v \in C \setminus D$. Because w is a neighbour of v and (C, D) is a separation, w is contained in C and thus for the corridor γ containing (C, D) we have that A_γ contains both v and w. □

Corollary 7.10 *Let F be a finite connected set of vertices not meeting Π. Then there is a separation $(C, D) \in O$ such that $F \subseteq C \setminus D$.*

Proof By Lemma 7.9 we may apply Lemma 7.6. □

Lemma 7.11 *Let γ be a corridor and assume that all separators of separations in N are cliques. Then $A_\gamma \cap \Pi$ is a clique, too.*

Proof Let v and w be two distinct vertices of $A_\gamma \cap \Pi$. Then by Lemma 7.6 there is a separation $(C, D) \in \gamma$ such that C contains both v and w. Because v and w are contained in Π which in turn is a subset of D, both v and w are contained in $C \cap D$. Because $C \cap D$ is a clique by assumption, v is a neighbour of w. □

8 Extending the tree set of the principal vertex sets

In this section we prove our main result, Theorem 1. To obtain a starting tree set T as described in our overall proof strategy in Sect. 4, we apply our technical main result Theorem 5.12 (combined with Theorem 5.10) to a carefully chosen collection $\mathcal{Y}$ of principal vertex sets of G. For choosing $\mathcal{Y}$ we need the following definition:

Definition 8.1 A separation $\{X, \mathscr{C}\}$ is *generous* if both $\mathscr{C}$ and the complement $\mathscr{C}_X \setminus \mathscr{C}$ contain components whose neighbourhoods are precisely equal to X, i.e. if $\breve{\mathscr{C}}_X$ meets both $\mathscr{C}$

and $\mathscr{C}_X \setminus \mathscr{C}$. A set X of vertices of G is *generous* if it is the separator of some generous separation, i.e. if $|\breve{\mathscr{C}}_X| \geq 2$.

Now we are ready to set up our starting tree set T and more, as follows. **Throughout this section we fix the following notation.** We let $\mathcal{Y}$ be the collection of all generous subsets of the critical vertex sets of G, in formula:

$$\mathcal{Y} = \left\{ X \in \mathcal{X} \mid X \text{ is generous and } \exists Y \in \mathrm{crit}(G) : X \subseteq Y \right\}.$$

Notably, $\mathrm{crit}(G) \subseteq \mathcal{Y}$. The elements of $\mathcal{Y}$ are principal (cf. Example 5.3). We assume, without loss of generality by Lemma 5.4, that each $X \in \mathcal{Y}$ induces a clique $G[X]$. Using Theorem 5.10, Remark 5.11 and Theorem 5.12 we obtain a strongly admissable function $\mathscr{K}$ for $\mathcal{Y}$ that deviates from all $\mathscr{C}_X$ with $X \in \mathcal{Y}$ by precisely one component, in formula $|\breve{\mathscr{C}}_X \setminus \mathscr{K}(X)| = 1$ for all $X \in \mathcal{Y}$. This way we ensure that $T := T(\mathcal{Y}, \mathscr{K})$ is a regular tree set of generous finite-order separations of G. For $X \in \mathcal{Y}$ we abbreviate $\sigma_X = \sigma_X^{\mathscr{K}}$. Moreover, O always denotes a consistent orientation of T, and then $\Pi \subseteq V(G)$ denotes the part of O. At some point in this section the concept of a 'modified torso' of O will be defined. From that point onward, H will always denote the modified torso of O. Whenever we speak of Π or H we tacitly assume that they stem from some O. This completes the list of fixed notation for this section.

Next, we consider two inequivalent $\aleph_0$-tangles τ_1 and τ_2 of G, we pick a finite-order separation $\{A_1, A_2\}$ of G that efficiently distinguishes τ_1 and τ_2, and we write $Z = A_1 \cap A_2$ for its separator. If Z is included entirely in a critical vertex set of G, then T efficiently distinguishes τ_1 and τ_2:

Lemma 8.2 *Let $\{Z, \mathscr{D}\}$ efficiently distinguish two $\aleph_0$-tangles τ_1 and τ_2 of G. Then $\{Z, \mathscr{D}\}$ is generous. If additionally τ_1 and τ_2 are inequivalent and Z is included in some critical vertex set of G, then T efficiently distinguishes τ_1 and τ_2.*

Proof Let $\{\mathscr{D}_1, \mathscr{D}_2\} := \{\mathscr{D}, \mathscr{C}_Y \setminus \mathscr{D}\}$ such that $(Z, \mathscr{D}_i) \in \tau_i$ for both $i = 1, 2$. Our proof starts with a more general analysis of the situation, as follows. Consider any $i \in \{1, 2\}$ and put $j = 3 - i$.

If τ_i lives in a component C of $G - Z$ in that $(Z, C) \in \tau_i$, then by the consistency of τ_j we deduce from $(C, Z) \leq (Z, \mathscr{D}_j) \in \tau_j$ that $(C, Z) \in \tau_j$, so $\{Z, C\}$ distinguishes τ_1 and τ_2. But then so does $\{N(C), C\}$ by Lemma 2.7, and hence $N(C) = Z$ follows by the efficiency of Z.

Otherwise τ_i is an ultrafilter tangle and $X_i := X_{\tau_i}$ is contained in Z. Then, as $U(\tau_i, Z)$ is a free ultrafilter, we have $(Z, \mathscr{D}_i') \in \tau_i$ for $\mathscr{D}_i' := \mathscr{D}_i \cap \mathscr{C}_Z(X_i)$. Hence $(X_i, \mathscr{D}_i') \in \tau_i$ by Lemma 2.7. And $(\mathscr{D}_i', X_i) \leq (Z, \mathscr{D}_j) \in \tau_j$ implies $(\mathscr{D}_i', X_i) \in \tau_j$ by the consistency of τ_j. Therefore, $\{X_i, \mathscr{D}_i'\}$ distinguishes τ_1 and τ_2, so $X_i = Z$ follows by the efficiency of Z.

From the two cases above we deduce that $\{Z, \mathscr{D}\}$ is generous. It remains to show that if additionally τ_1 and τ_2 are inequivalent and Z is contained in a critical vertex set of G, then T efficiently distinguishes τ_1 and τ_2. First, we have $Z \in \mathcal{Y}$ as Z is generous. Next, we note that not both τ_1 and τ_2 can be ultrafilter tangles with $X_1, X_2 \subseteq Z$ for otherwise $X_1 = Z = X_2$ follows from our considerations above, contradicting that τ_1 and τ_2 are inequivalent. So at least one of τ_1 and τ_2 lives in a component C of $G - Z$, say $(Z, C) \in \tau_1$, and then $C \in \breve{\mathscr{C}}_Z$ follows from our considerations above. If $C \in \mathscr{K}(Z)$ then $\{Z, C\} \in T$ efficiently distinguishes τ_1 and τ_2. Otherwise $\{C\} = \breve{\mathscr{C}}_Z \setminus \mathscr{K}(Z)$, and we claim that $\{Z, \mathscr{K}(Z)\} \in T$ efficiently distinguishes τ_1 and τ_2. On the one hand, $(\mathscr{K}(Z), Z) \leq (Z, C) \in \tau_1$ implies

$(\mathscr{K}(Z), Z) \in \tau_1$ by the consistency of τ_1. On the other hand, τ_2 either lives in a component in $\mathscr{K}(Z)$ or τ_2 is an ultrafilter tangle with $X_2 = Z$, and in both cases we deduce $(Z, \mathscr{K}(Z)) \in \tau_2$. $\qquad\square$

Therefore, we may assume that Z is not contained entirely in any critical vertex set of G. Then Z is contained in a part of T, as follows.

Lemma 8.3 *Let $Z \in \mathcal{X}$ be generous. If X is a principal vertex set of G that does not contain Z entirely, then there is a unique component of $G - X$ that Z meets.*

Proof As Z is not contained in X as a subset, there is a component C of $G - X$ which Z meets. Assume for a contradiction that there is another component D of $G - X$ meeting Z. Pick vertices $c \in Z \cap C$ and $d \in Z \cap D$. Now note that every component $K \in \check{\mathscr{C}}_Z$ must meet X, for K plus its K–c and K–d edges admits a c–d path connecting the distinct components C and D of $G - X$. But since X is principal it meets at most one component of $G - Z$, namely $C_Z(X)$, contradicting that $|\check{\mathscr{C}}_Z| \geq 2$. $\qquad\square$

By Lemma 8.3 above the separator Z meets precisely one side from every separation in T, and then orienting each separation in T towards that side results in a consistent orientation O of T whose part Π contains Z. The remainder of this section is dedicated to modifying the torso H of O so that

- τ_1 and τ_2 are 'represented' by 'proxy' ends η_1 and η_2 in H; and
- applying Carmesin's theorem in H yields a tree set that lifts compatibly with T to a tree set of tame finite-order separations of G that efficiently distinguishes τ_1 and τ_2.

8.1 Modified torsos, proxies of corridors and lifting from modified torsos

In this subsection we introduce modified torsos and show that there is an elegant way to lift tree sets from modified torsos to the graph G itself. Proxies of corridors are introduced as a technical tool whose purpose is twofold: first, they are key to the elegant lifting of tree sets. And second, they will be employed in the next subsection to define proxies for $\aleph_0$-tangles.

Definition 8.4 (Modified torso) Whenever Π is non-empty we define the *modified torso H* of O, as follows. Consider the set $\mathcal{Z}$ of all finite subsets of Π that are separators of suprema of corridors of O. Then we obtain H from $G[\Pi]$ by disjointly adding for each $X \in \mathcal{Z}$ a copy of $K^{\aleph_0}$ that we join completely to X.

We remark that Π being non-empty ensures that the empty set is not an element of $\mathcal{Z}$, so modified torsos are connected. Since the copies of $K^{\aleph_0}$ are joined to finite cliques of $G[\Pi]$ by Lemma 7.11, no two ends of $G[\Pi]$ are merged when we move on to the modified torso H:

Lemma 8.5 *Every finite-order separation of $G[\Pi]$ extends to a finite-order separation of H. Thus sending each end η of $G[\Pi]$ to the end $\iota(\eta)$ of $H \supseteq G[\Pi]$ with $\eta \subseteq \iota(\eta)$ defines an injection $\iota : \Omega(G[\Pi]) \hookrightarrow \Omega(H)$. Moreover, the ends of H that do not lie in the image of ι correspond bijectively to the copies of $K^{\aleph_0}$ that were added to $G[\Pi]$ in order to obtain H.* $\square$

Now we tend to the lifting of separations from H to G. It is desirable to have the separator of a separation remain unchanged when lifting it. But H usually will contain many vertices that are not vertices of the original graph G. We solve this as follows. When we consider finite-order separations of H, we are only interested in ones that efficiently distinguish some two ends of H. And these H-relevant separations have their separators consist of vertices of the original graph G:

Definition 8.6 (*H-relevant*) If a separation of H has finite order and efficiently distinguishes some two ends of H, then we call it and its orientations *H-relevant*.

Lemma 8.7 *If $\{A, B\}$ is H-relevant, then $A \cap B \subseteq \Pi$.*

Proof Assume for a contradiction that $A \cap B$ meets an added copy K of a $K^{\aleph_0}$ in a vertex v and write $X = N_H(K)$. Notably $H[X]$ is a clique, and hence so is $H[X \cup K]$. Without loss of generality we may assume that $H[X \cup K] \subseteq H[A]$, so K meets $A \setminus B$ while $H[X \cup K]$ avoids $B \setminus A$. Now $v \in A \cap B \cap K$ sends its edges only to K and X, and in particular v sends no edges to $B \setminus A$. So $\{A, B - v\}$ is again a separation of H, but of order $|A \cap B| - 1$, and this separation still distinguishes all the ends of H that were distinguished by $\{A, B\}$, contradicting that $\{A, B\}$ is H-relevant. □

Now we are almost ready to define lifts of separations of H, all we are missing is

Definition 8.8 (*Proxy of a corridor*) Suppose that Π is non-empty and γ is a corridor of O. The *proxy* of γ in the modified torso H is the end η of H that is defined as follows. Consider the separator X of the supremum of γ. If X is finite, then η is the end of H containing the rays of the $K^{\aleph_0}$ that was added for X. Otherwise $G[X] \subseteq H$ is an infinite clique by Lemma 7.11, and then η is the end of H that contains the rays of $G[X]$.

Finally, we can lift separations from H to G:

Definition 8.9 (*Lift from a modified torso*) Let (A, B) be an H-relevant separation of a modified torso H. By Lemma 8.7 the separator $A \cap B$ is included in Π entirely. The *lift* $(\ell(A), \ell(B))$ of (A, B) is defined as follows. The set $\ell(A) \subseteq V(G)$ agrees with A on Π, and a vertex of $G - \Pi$ is contained in $\ell(A)$ whenever its corridor's proxy in H lives on the A-side. The set $\ell(B)$ is defined analogously.

We remark that $\{\ell(A), \ell(B)\}$ does not depend on the orientation of $\{A, B\}$. In order to verify that the lifts work as intended we need the following lemma:

Lemma 8.10 *If $\{A, B\}$ is H-relevant and γ is a corridor of O whose proxy lives on the A-side, then $A_\gamma \subseteq \ell(A)$.*

Proof We have $A_\gamma \setminus \Pi \subseteq \ell(A)$ by the definition of $\ell(A)$. It remains to show $\Xi \subseteq A$ for the separator $\Xi = A_\gamma \cap \Pi$ of the supremum of γ. If Ξ is infinite, then the proxy η of γ living on the A-side means $G[\Xi] \subseteq H[A]$. Otherwise Ξ is finite. Then η stems from a copy $K \subseteq H$ of $K^{\aleph_0}$ that is joined completely to the clique $G[\Xi]$, and so η living on the A-side means $K \subseteq H[A]$. Consequently, the infinite clique $H[K \cup \Xi]$ is contained in $H[A]$ as well, giving $\Xi \subseteq A$ as desired. □

Now we can check for ourselves that lifts work:

Lemma 8.11 *The lift of an H-relevant separation is a separation of G with the same separator.*

Proof Let $\{A, B\}$ be any H-relevant separation, and recall that $A \cap B \subseteq \Pi$ by Lemma 8.7. Every vertex of $G - \Pi$ lies in A_γ for a unique corridor γ of O, and hence is contained in precisely one of $\ell(A)$ and $\ell(B)$. Thus $\ell(A) \cap \ell(B) = A \cap B$. It remains to verify that G has no edge between $\ell(A) \setminus \ell(B)$ and $\ell(B) \setminus \ell(A)$. For this, let $e = xy$ be any edge of G. If both x and y are contained in Π, then $e \subseteq A$ say, and hence $e \subseteq \ell(A)$. Otherwise one of x and y lies outside of Π, say $x \in \ell(A) \setminus \Pi$. Let γ be the corridor of O with $x \in A_\gamma \setminus \Pi$, so the proxy η of γ lives on the A-side. From $x \in A_\gamma \setminus \Pi$ we infer $y \in A_\gamma$. Then $e \subseteq A_\gamma \subseteq \ell(A)$ by Lemma 8.10. □

Starting with an intuitive lemma we verify that our lifts are compatible with T and lifts of other modified torsos:

Lemma 8.12 *Let γ be a corridor of O and let η be the proxy of γ in H. If $\{A, B\}$ is H-relevant with η living on the A-side, then $\vec{S} \leq (\ell(A), \ell(B))$ for all $\vec{S} \in \gamma$. In particular, the lift of an H-relevant separation is nested with T.*

Proof Consider any $(C, D) \in \gamma$. We have to show $(C, D) \leq (\ell(A), \ell(B))$. For the inclusion $C \subseteq \ell(A)$ we start with $C \subseteq A_\gamma$ and employ Lemma 8.10 for $A_\gamma \subseteq \ell(A)$. Now the inclusion $\ell(B) \subseteq D$ is tantamount to $C \setminus D \subseteq \ell(A) \setminus \ell(B)$ which is immediate from $C \subseteq \ell(A)$ as $C \setminus D$ avoids $\Pi \supseteq \ell(A) \cap \ell(B)$ (cf. Lemma 8.11). □

Corollary 8.13 *If H' is the modified torso of a consistent orientation O' of T other than O, then all lifts of H-relevant separations are nested with all lifts of H'-relevant separations.*

Lemma 8.14 *If (A, B) and (C, D) are two H-relevant with $(A, B) \leq (C, D)$, then their lifts satisfy $(\ell(A), \ell(B)) \leq (\ell(C), \ell(D))$. In particular, the lifts of two nested H-relevant separations are again nested.*

We close this subsection with the lemma that ensures that when we construct the tree set for our main result, we are able to ensure the 'moreover' part stating equivalent $\aleph_0$-tangles orient the tree set the same way.

Lemma 8.15 *Every H-relevant separation lifts to a tame separation of G.*

Proof If H stems from a consistent orientation O of T that contains some star σ_X with $X \in \mathcal{Y}$, then $O = \lceil \sigma_X \rceil$ (with the down-closure taken in $\vec{T}$) by Theorem 5.12. Consequently, H was obtained from the finite clique $G[X]$ by disjointly adding precisely one copy of a $K^{\aleph_0}$ and joining it completely to X, so the one-ended H has no H-relevant separations. Therefore, we may assume that O avoids all of the stars σ_X with $X \in \mathcal{Y}$.

Let $\{A, B\}$ be an H-relevant separation and recall that $A \cap B \subseteq \Pi$ by Lemma 8.7. And let $X \subseteq A \cap B$ be a critical vertex set of G.

If there is a component $K \in \mathscr{K}(X)$ with $(X, K) \in O$, then the proxy of the corridor of O that contains (X, K) ensures that all the components in $\mathscr{K}(X) \setminus \{K\}$ are contained in the same side of $\{\ell(A), \ell(B)\}$.

Otherwise, since O avoids the star σ_X, we have $(\mathscr{K}(X), X) \in O$. Then the proxy of the corridor of O that contains $(\mathscr{K}(X), X)$ ensures that all the components in $\mathscr{K}(X)$ are contained in the same side of $\{\ell(A), \ell(B)\}$.

In either case, all but finitely many of the components in $\mathscr{C}_X$ lie on the same side of $\{\ell(A), \ell(B)\}$. Since $A \cap B \supseteq X$ meets at most finitely many components in $\mathscr{C}_X$, the collection $\mathscr{C}_{A \cap B}(X)$ forms a cofinite subset of $\check{\mathscr{C}}_X$, and therefore all but finitely many components in $\mathscr{C}_{A \cap B}(X)$ lie on the same side of $\{\ell(A), \ell(B)\}$ as desired. $\qquad \square$

8.2 Proxies of $\aleph_0$-tangles

We start this subsection by introducing the technical notion of 'walking a corridor' and prove two technical lemmas about ends. This framework, together with proxies of corridors, then enables us to give a comprehensible definition of proxies of $\aleph_0$-tangles. We emphasise that this technical layering is highly important to save the key segments of our overall proof from being swamped with terrible amounts of case distinctions.

Definition 8.16 (Walking) We say that an end ω of G *walks* a corridor γ of O if for the supremum (A, B) of γ the end ω has a ray contained in $G[A \setminus B]$. And we say that an ultrafilter tangle τ of G *walks* a corridor γ of O if τ contains the inverse of some separation in γ.

Lemma 8.17 *Suppose that N is a tree set of generous finite-order separations of G all whose separators induce cliques. Let ω be an end of G, let Π be the part of the orientation $O = \omega \cap \vec{N}$ that ω induces on N, and suppose that ω walks a corridor γ of O. If the separator $A_\gamma \cap \Pi$ of the supremum of γ is infinite, then $G[\Pi]$ contains a ray from ω.*

Proof Pick $R \in \omega$ arbitrarily. By Lemma 7.11 it is sufficient to show that there are infinitely many pairwise disjoint paths from R to $A_\gamma \cap \Pi$. We will recursively construct such paths P_n ($n \in \mathbb{N}$) of which only the last vertex v_n is contained in Π. Assume that $P_0, \ldots, P_{n-1}$ have already been defined. Then there is a finite non-empty initial segment R' of R such that $R' \cup P_0 \mathring{v}_0 \cup \cdots \cup P_{n-1} \mathring{v}_{n-1}$ is connected. Let $(A, B) \in O$ be a separation such that all vertices of $R' \cup P_0 \mathring{v}_0 \cup \cdots \cup P_{n-1} \mathring{v}_{n-1}$ are contained in $A \setminus B$ (such a separation exists by Corollary 7.10). Then (A, B) is contained in γ. Every vertex v_k with $k < n$ is a neighbour of a vertex in $A \setminus B$ and thus contained in A.

As $A_\gamma \cap \Pi$ is infinite, it contains a vertex v which is not contained in $A \cap B$ and thus not contained in A. In particular, v is not contained in any path P_k with $k < n$. Because $v \in A_\gamma$, there is a separation (A', B') in γ such that $v \in A'$ and thus $v \in A' \cap B'$. Let (C, D) be a separation in γ which is bigger than both (A, B) and (A', B'). Then all vertices contained in some P_k with $k < n$ are contained in $(C \setminus D) \cup \Pi$. Furthermore $v \in C \cap D$ and R contains a vertex of $C \setminus D$. Because $(C, D) \in O$, some tail of R is contained in $D \setminus C$ and thus R also contains vertices of $D \setminus C$. As (C, D) is a separation and R connected this implies that some vertex w of R is contained in $C \cap D$. Because w is not a vertex of Π it is also not a vertex of some P_k with $k < n$.

As (D, C) is generous, there is a component of $G - (C \cap D)$ which is contained in $D \setminus C$ and whose neighbourhood is precisely equal to $C \cap D$. Thus there is a path P from w to v whose inner vertices are contained in $D \setminus C$. We already established that v and w are not vertices of any P_k with $k < n$. Hence P is disjoint from all P_k with $k < n$. Let v_n be the first vertex of P in Π and let $P_n := wPv_n$. By Corollary 7.10 there is a separation $(I, J) \in O$ such that the vertices of $P_n \mathring{v}_n$ are contained in $I \setminus J$. Then $(I, J) \in \gamma$ and $v_n \in I$, so $v_n \in A_\gamma$. As also $v_n \in \Pi$ we have $v_n \in A_\gamma \cap \Pi$ as required.   □

Lemma 8.18 *If an end ω of G does not lie in the closure of Π, then ω walks a unique corridor of O.*

Proof   Since ω does not lie in the closure of Π, we in particular find a ray $R \in \omega$ that avoids Π. As R is connected, it defines a corridor γ of O with $R \subseteq A_\gamma \setminus \Pi$. Then ω walks the corridor γ, and so it remains to show that γ is unique.

If $O \not\subseteq \omega$, then ω contains the inverse $\overleftarrow{s}$ of some separation $\overrightarrow{S} \in O$, and then γ is determined as the corridor of O containing $\overrightarrow{S}$. Otherwise $O \subseteq \omega$. Then we assume for a contradiction that there is another ray $R' \in \omega$ that walks a corridor γ' of O other than γ. Since the suprema of γ and γ' both separate R and R', their separators cannot be finite, and so they are infinite. But then applying Lemma 8.17 to either γ or γ' yields a ray of ω in $G[\Pi]$, contradicting the assumption that ω does not lie in the closure of Π.   □

Finally, we are ready for the definition of proxies of $\aleph_0$-tangles. We split the definition and consider ends and ultrafilter tangles separately.

Definition 8.19   (Proxy of an end) If ω is an end of G, then the *proxy* of ω in H is the end η of H that is defined as follows.

- If ω lies in the closure of Π, then ω has a ray in $G[\Pi]$ by Corollary 2.4, and η is the end of such a ray in H (this is well-defined by Corollary 2.5).
- Otherwise ω does not lie in the closure of Π and by Lemma 8.18 walks a unique corridor γ of O; then η is the proxy of γ in H.

Definition 8.20   (Proxy of an ultrafilter tangle) If τ is an ultrafilter tangle of G and O avoids the star σ_{X_τ}, then τ walks a unique corridor γ of O and the *proxy* of τ in H is the end η of H that is the proxy of γ in H.

We close this subsection with a lemma on the interaction of lifts with proxies:

Lemma 8.21   *Let τ be an $\aleph_0$-tangle of G and suppose that the proxy η of τ in H is defined. If $\{A, B\}$ is H-relevant and $(A, B) \in \eta$, then $(\ell(A), \ell(B)) \in \tau$.*

Proof   If τ is an end of G that lies in the closure of Π, then this follows from the fact that some ray of $G[\Pi]$ is contained in both τ and η. Otherwise τ is an $\aleph_0$-tangle of G that walks a unique corridor γ of O. If additionally τ is an end, then every ray in τ that avoids Π is contained in $\ell(B)$, ensuring $(\ell(A), \ell(B)) \in \tau$. So we may assume that τ is an ultrafilter tangle. As the proxy η of τ is defined, we know that O avoids the star σ_{X_τ} so that τ walks a unique corridor γ of O. By definition, this means that τ contains the inverse of some oriented

separation from γ. Then $(\ell(A), \ell(B)) \in \tau$ follows from Lemma 8.12 and the consistency of the tangle τ. $\qquad \square$

8.3 Efficiently distinguishing the proxies

In this subsection we provide the final key segments of our overall proof. We start with an overview of the situation that is of interest.

Throughout this subsection we fix the following notation in addition to the notation fixed throughout the ambient section. (See also Fig. 7.) We are given two inequivalent $\aleph_0$-tangles τ_1 and τ_2 of G that are efficiently distinguished by a finite-order separation $\{A_1, A_2\}$ of G with separator $Z = A_1 \cap A_2$. The separator Z is not contained in a critical vertex set of G. Hence, by Lemma 8.3 the separator Z meets precisely one side from every separation in T, and then orienting each separation in T towards that side results in a consistent orientation O of T whose part Π contains Z. For this special orientation we write $O(Z)$, and we write $\Pi(Z)$ and $H(Z)$ for its part and modified torso. Moreover, η_1 and η_2 are the proxies of τ_1 and τ_2 in $H(Z)$ (note that these are defined as $O(Z)$ avoids all stars σ_X with $X \in \mathcal{Y}$). Whenever we write i we mean an arbitrary $i \in \{1, 2\}$, and we write $j = 3 - i$. If τ_i happens to be an ultrafilter tangle, then we write X_i instead of X_{τ_i}. This completes the list of fixed notation for this subsection.

The final key segments are Lemma 8.22 and Proposition 8.24 below. We start with the lemma:

Lemma 8.22 *Every relevant finite-order separation of H that distinguishes η_1 and η_2 does lift to a separation of G that distinguishes τ_1 and τ_2.*

Proof Let $\{A, B\}$ be any relevant finite-order separation of H that distinguishes η_1 and η_2, say with $(B, A) \in \eta_1$ and $(A, B) \in \eta_2$. Then Lemma 8.21 gives both $(\ell(B), \ell(A)) \in \tau_1$ and $(\ell(A), \ell(B)) \in \tau_2$, so $\{\ell(A), \ell(B)\}$ distinguishes τ_1 and τ_2. $\qquad \square$

For the key proposition, we need the following proposition whose proof we postpone to after the proof of the key proposition.

Proposition 8.23 *If τ_i walks a corridor γ_i of $O(Z)$ where Ξ_i denotes the separator of the supremum of γ_i, then $G[\Xi_i \setminus Z]$ is a non-empty clique that is entirely contained in $G[A_i \setminus A_j]$.*
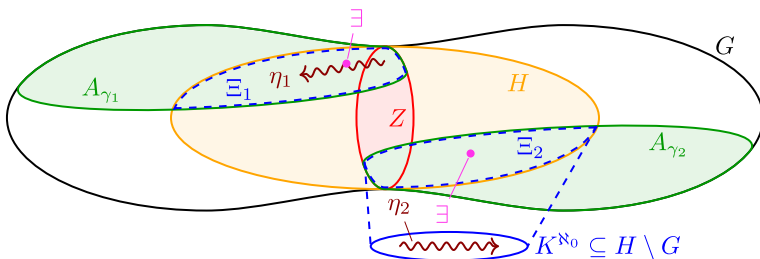
The final key segment is

Fig. 7 A graph G, a vertex set Z efficiently separating two inequivalent $\aleph_0$-tangles τ_1 and τ_2 of G, and a modified torso $H(Z)$ which contains Z. The $\aleph_0$-tangles τ_i walk corridors γ_i and $H(Z)$ has proxies η_1 and η_2 for τ_1 and τ_2

Proposition 8.24 *The proxies η_1 and η_2 are efficiently distinguished by Z.*

Proof If Z distinguishes the proxies η_1 and η_2 in $H(Z)$, then it does so efficiently, for otherwise the separation of order $< |Z|$ doing so lifts to one distinguishing τ_1 and τ_2 in G by Lemma 8.22, contradicting the efficiency of Z. Therefore, it remains to show that η_1 and η_2 are distinguished by Z. For this, we check three cases.

In the first case, both τ_1 and τ_2 lie in the closure of $\Pi(Z)$. Then τ_1 and τ_2 are distinct ends of G that lie in the closure of $\Pi(Z)$, and so their proxies stem from rays of τ_1 and τ_2 respectively. Now Z witnesses that these rays are inequivalent in G and, in particular, that they are inequivalent in $G[\Pi]$. Thus Z distinguishes η_1 and η_2 in $H(Z)$ by Lemma 8.5.

In the second case, neither τ_1 nor τ_2 lies in the closure of $\Pi(Z)$, and both walk corridors γ_1 and γ_2 of $O(Z)$. We let Ξ_1 and Ξ_2 be the separators of the suprema of γ_1 and γ_2. Then, by Proposition 8.23, for both $i = 1, 2$ the induced subgraph $G[\Xi_i \setminus Z]$ is a non-empty clique that is entirely contained in $G[A_i \setminus A_j]$. Consequently, Z distinguishes η_1 and η_2 in $H(Z)$ by Lemma 8.5.

In the third case, τ_1 does not lie in the closure of $\Pi(Z)$ and walks a corridor γ_1 of $O(Z)$ while τ_2 lies in the closure of $\Pi(Z)$. Then τ_2 must be an end of G. We let Ξ_1 be the separator of the supremum of γ_1.

By Proposition 8.23 the induced subgraph $G[\Xi_1 \setminus Z]$ is a non-empty clique that is entirely contained in $G[A_1 \setminus A_2]$. Since η_1 stems from the copy of $K^{\aleph_0}$ that is attached to the clique $G[\Xi_1] \subseteq G[\Pi]$ while η_2 stems from a ray of $G[\Pi]$ in τ_2, we deduce that Z distinguishes η_1 and η_2 in $H(Z)$ by Lemma 8.5. □

In the remainder of this subsection we prove Proposition 8.23. For this, we introduce the concept of a pointer. Basically, the idea is to have a connected subgraph of G that can be employed as an oracle—like we employ rays as oracles for their ends.

Definition 8.25 (Pointer) If τ_i walks a corridor γ_i of $O(Z)$, then a *pointer* of τ_i is a connected subgraph K_i of $G[A_{\gamma_i} \setminus \Pi(Z)] \cap G[A_i \setminus A_j]$ that is of the following form. If τ_i is an end of G, then K_i is a ray in τ_i. Otherwise τ_i is an ultrafilter tangle of G, and then K_i is a component in $\breve{\mathscr{C}}_{X_i}$.

Lemma 8.26 *If τ_i walks a corridor of $O(Z)$, then τ_i has a pointer.*

Proof If τ_i is an end, then τ_i has a ray avoiding $\Pi(Z)$ for τ_i walks a corridor of $O(Z)$, and as $Z \subseteq \Pi(Z)$ every such ray is a pointer of τ_i. Otherwise τ_i is an ultrafilter tangle. Then we let γ_i be the corridor of $O(Z)$ walked by τ_i. Let $(C, D) \in \tau_i$ witness that τ_i walks γ_i, so $D \setminus C \subseteq A_{\gamma_i} \setminus \Pi(Z)$. Using Theorem 2.8 we pick $\mathscr{C} \in U(\tau_i, X_i)$ with $V[\mathscr{C}] \subseteq A_i \setminus A_j$ and $\mathscr{C}' \in U(\tau_i, X_i)$ with $V[\mathscr{C}'] \subseteq D \setminus C$. As $U(\tau_i, X_i)$ is a free ultrafilter, the intersection $\mathscr{C} \cap \mathscr{C}' \cap \breve{\mathscr{C}}_{X_i} \in U(\tau_i, X_i)$ is infinite, and every component in this intersection is a pointer of τ_i. □

Lemma 8.27 *If the neighbourhood $N(C_i)$ of the component C_i of $G - \Pi(Z)$ containing a pointer K_i of τ_i is finite, then $N(C_i)$ is the separator of some finite-order separation of G that distinguishes τ_1 and τ_2.*

Proof By the consistency of τ_j it suffices to find a separation $(A, B) \in \tau_i$ with $(B, A) \leq (A_i, A_j)$ and $A \cap B = Y$ where we write $Y = N(C_i)$. As K_i is a pointer we have

$K_i \subseteq G[A_i \setminus A_j]$. Since the separator $Z = A_1 \cap A_2$ is included in $\Pi(Z)$ we have $C_i \subseteq G[A_i \setminus A_j]$ as well. If τ_i is an end then $(Y, C_i) \in \tau_i$ is as desired. Otherwise τ_i is an ultrafilter tangle. If $C_i = K_i$ then $N(C_i) = N(K_i) = X_i$, and employing Theorem 2.8 we may pick $\mathscr{C} \in U(\tau_i, X_i)$ with $V[\mathscr{C}] \subseteq A_i$, so $(X_i, \mathscr{C}) \in \tau_i$ is as desired. Hence we may assume that $C_i \supsetneq K_i$ must meet X_i. Then $C_i = C_Y(X_i)$, and by Lemma 2.10 we have $(Y, C_Y(X_i)) \in \tau_i$ as desired. □

Proof of Proposition 8.23 By Lemma 8.26 we find a pointer K_i of τ_i, and we let C_i be the component of $G - \Pi(Z)$ containing K_i. Then $K_i \subseteq G[A_{\gamma_i} \setminus \Pi(Z)]$ implies $C_i \subseteq G[A_{\gamma_i} \setminus \Pi(Z)]$, so we have $N(C_i) \subseteq \Xi_i$. First, we show that C_i has a neighbour in $\Xi_i \setminus Z$. Otherwise $N(C_i) \subseteq \Xi_i \cap Z$, and then $N(C_i) = Z$ by Lemma 8.27 and the efficiency of Z. Now $Z \subseteq \Xi_i$ with Z being finite allows us to find a separation $(X, \mathscr{C}) \in \gamma_i$ with $Z \subseteq X$ contradicting this subsection's assumption on Z. Therefore, C_i has a neighbour in $\Xi_i \setminus Z$. Next, since $G[\Xi_i]$ is a clique, there is a unique component D_i of $G - Z$ containing $G[\Xi_i \setminus Z]$. Then $C_i \subseteq D_i$ as C_i has a neighbour in $\Xi_i \setminus Z$, and so $G[\Xi_i \setminus Z] \subseteq G[A_i \setminus A_j]$ follows from the pointer K_i being included in $G[A_i \setminus A_j]$. □

8.4 Proof of the main result

At last, we prove our main result:

Theorem 1 *Every connected graph G has a tree set of tame finite-order separations that efficiently distinguishes all its inequivalent $\aleph_0$-tangles. In particular, equivalent $\aleph_0$-tangles induce the same orientations on the tree set.*

Proof For every modified torso H of T we employ Carmesin's Theorem 2.13 to obtain a tree set T_H of H-relevant separations that efficiently distinguishes all the ends of H. Then we lift all the separations in all the tree sets T_H and add these to T to obtain an extension T' of T. Then T' is again a tree set by Lemma 8.12, Corollary 8.13 and Lemma 8.14.

First, we show that T' efficiently distinguishes every two inequivalent $\aleph_0$-tangles of G. For this, let τ_1 and τ_2 be two inequivalent $\aleph_0$-tangles of G. We have to find a separation in T' that efficiently distinguishes τ_1 and τ_2. Pick some finite-order separation of G with separator Z say that efficiently distinguishes τ_1 and τ_2. If Z is contained in some critical vertex set of G, then by Lemma 8.2 we find a separation in $T \subseteq T'$ that efficiently distinguishes τ_1 and τ_2. Otherwise Z is not contained in any critical vertex set of G. However, Z is generous by Lemma 8.2, and so by Lemma 8.3 induces a consistent orientation of T whose part contains Z. Then by Proposition 8.24 the $\aleph_0$-tangles τ_1 and τ_2 have distinct proxies η_1 and η_2 in the modified torso H of that orientation, and Z efficiently distinguishes η_1 and η_2 in H. Thus there is a separation in T_H of order $|Z|$ that distinguishes η_1 and η_2. By Lemma 8.22 this separation lifts to a separation of G that distinguishes τ_1 and τ_2. This lift still has order $|Z|$ and lies in T'.

Second we show that all separations in T' are tame. Every separation in T is tame. And by Lemma 8.15 the lifts of all T_H are tame as well. $\square$

Appendix

Compactifications

A *compactification* of a topological space X is an ordered pair (K, h) where K is a compact topological space and $h : X \hookrightarrow K$ is an embedding of X as a dense subset of K. Sometimes we also refer to K as a compactification of X if the embedding h is clearly understood (e.g. if h is the identity on X). The space $K \setminus h[X]$ is called the *remainder* of the compactification.

Suppose now that X is a discrete topological space. Since X is locally compact[5], X is open in all of its Hausdorff compactifications (cf. [24, Theorem 3.6.6]). If we pair the space βX of all ultrafilters on X carrying the topology whose basic open sets are of the form $\{U \in \beta X \mid A \in U\}$, one for each $A \subseteq X$, with the embedding that sends every $x \in X$ to the principal ultrafilter on X generated by $\{x\}$, then this yields the Stone-Čech compactification of X. By the Stone-Čech property, every continuous function $f : X \to T$ into a compact Hausdorff space T has a unique continuous extension $\beta f : \beta X \to T$ with $\beta f \mid X = f$ (cf. [24, Theorem 3.5.1]).

Theorem 9.1 ([26, Corollary 7.4]) *If X is an infinite set, then $|\beta X| = 2^{2^{|X|}}$.*

Inverse limits of inverse systems

A partially ordered set $(I, \leq)$ is said to be *directed* if for every two $i, j \in I$ there is some $k \in I$ with $k \geq i, j$. Let $(X_i \mid i \in I)$ be a family of topological spaces indexed by some directed poset $(I, \leq)$. Furthermore, suppose that we have a family $(\varphi_{ji} : X_j \to X_i)_{i \leq j \in I}$ of continuous mappings which are the identity on X_i in case of $i = j$ and which are *compatible* in that $\varphi_{ki} = \varphi_{ji} \circ \varphi_{kj}$ for all $i \leq j \leq k$. Then both families together are said to form an *inverse system*, and the maps φ_{ji} are called its *bonding maps*. We denote such an inverse system by $\{X_i, \varphi_{ji}, I\}$ or $\{X_i, \varphi_{ji}\}$ for short if I is clear from context. Its *inverse limit* $\varprojlim X_i = \varprojlim(X_i \mid i \in I)$ is the topological space

$$\varprojlim X_i = \{(x_i)_{i \in I} \mid \varphi_{ji}(x_j) = x_i \text{ for all } i \leq j\} \subseteq \prod_{i \in I} X_i.$$

Whenever we define an inverse system without specifying a topology for the spaces X_i first, we tacitly assume them to carry the discrete topology. If each X_i is (non-empty) compact Hausdorff, then so is the inverse limit $\varprojlim X_i$.

---

[5] A topological space X is *locally compact* if for each of its points there is some compact subspace of X which includes an open neighbourhood of that point.

References

1. Diestel, R.: Graph Theory, 5th edn. Springer, Berlin, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53622-3
2. Robertson, N., Seymour, P.D.: Graph minors. X. obstructions to tree-decomposition. J. Combin. Theory Ser. B **52**, 153–190 (1991). https://doi.org/10.1016/0095-8956(91)90061-N
3. Carmesin, J.: All graphs have tree-decompositions displaying their topological ends. Combinatorica **39** (3), 545–596 (2019). https://doi.org/10.1007/s00493-018-3572-0. arXiv:1409.6640v5
4. Diestel, R.: Ends and Tangles. Abh. Math. Sem. Univ. Hamburg **87**(2), 223–244 (2017). https://doi.org/10.1007/s12188-016-0163-0. Special issue in memory of Rudolf Halin, arXiv:1510.04050v3
5. Diestel, R., Kühn, D.: Graph-theoretical versus topological ends of graphs. J. Combin. Theory Ser. B **87**, 197–206 https://doi.org/10.1016/S0095-8956(02)00034-5
6. Kurkofka, J., Pitz, M.: Ends, tangles and critical vertex sets. Math. Nachr. **292**(9), 2072–2091 (2019). https://doi.org/10.1002/mana.201800174. arXiv:1804.00588
7. Freudenthal, H.: Neuaufbau der endentheorie. Annals of Mathematics, pp. 261–279 (1942). https://doi.org/10.2307/1968869
8. Kurkofka, J., Pitz, M.: Tangles and the Stone-Čech compactification of infinite graphs. J. Combin. Theory Ser. B **146**, 34–60 (2021). https://doi.org/10.1016/j.jctb.2020.07.004. arXiv:1806.00220
9. Elbracht, C., Kneip, J., Teegen, M.: Trees of tangles in infinite separation systems. Mathematical Proceedings of the Cambridge Philosophical Society (2021). https://doi.org/10.1017/S0305004121000530. arxiv:2005.12122
10. Bürger, C., Kurkofka, J.: Duality theorems for stars and combs IV: Undominating stars. J. Graph Theory **100**, 140–162 (2022). https://doi.org/10.1002/jgt.22769. arXiv:2004.00591
11. Diestel, R.: End spaces and spanning trees. J. Combin. Theory Ser. B **96**, 846–854 https://doi.org/10.1016/j.jctb.2006.02.010
12. Polat, N.: Ends and multi-endings I. J. Combin. Theory (Series B) **67**, 86–110 (1996). https://doi.org/10.1006/jctb.1996.0035
13. Sprüssel, P.: End spaces of graphs are normal. J. Combin. Theory Ser. B **98**, 798–804 (2008). https://doi.org/10.1016/j.jctb.2007.10.006
14. Halin, R.: Über unendliche Wege in Graphen. Math. Annalen **157**, 125–137 (1964). https://doi.org/10.1007/BF01362670
15. Bürger, C., Kurkofka, J.: Duality theorems for stars and combs I: arbitrary stars and combs. J. Graph Theory **99**, 525–554 (2021). https://doi.org/10.1002/jgt.22706. arXiv:2004.00594
16. Bürger, C., Kurkofka, J.: Duality theorems for stars and combs II: dominating stars and dominated combs. J. Graph Theory **99**, 555–572 (2021). https://doi.org/10.1002/jgt.22757. arXiv:2004.00593
17. Bürger, C., Kurkofka, J.: Duality theorems for stars and combs III: undominated combs. J. Graph Theory **100**, 127–139 (2021). https://doi.org/10.1002/jgt.22781. arXiv:2004.00592
18. Diestel, R.: Abstract separation systems. Order **35**(1), 157–170 (2018). https://doi.org/10.1007/s11083-017-9424-5. arXiv:1406.3797v6
19. Diestel, R.: Tree sets. Order **35**(1), 171–192 (2018). https://doi.org/10.1007/s11083-017-9425-4. arXiv:1512.03781v3
20. Kneip, J., Gollin, P.: Representations of infinite tree sets. Order (2020). https://doi.org/10.1007/s11083-020-09529-0. arXiv:1908.10327

21. Carmesin, J., Hamann, M., Miraftab, B.: Canonical trees of tree-decompositions. J. Combin. Theory Ser. B **152**, 1–26 (2022). https://doi.org/10.1016/j.jctb.2021.08.004. arXiv:2002.12030
22. Dunwoody, M.J., Krön, B.: Vertex cuts. J. Graph Theory **80**(2), 136–171 (2015). https://doi.org/10.1002/jgt.21844
23. Bowler, N.: Hamburg Workshop on Abstract Separation Systems, Spiekeroog (2014)
24. Engelking, R.: General Topology, 2nd edn. Sigma Series in Pure Mathematics, vol. 6. Heldermann Verlag, Berlin (1989)
25. Kurkofka, J., Melcher, R., Pitz, M.: Approximating infinite graphs by normal trees. J. Combin. Theory Ser. B **148**, 173–183 (2021). https://doi.org/10.1016/j.jctb.2020.12.007. arXiv:2002.08340
26. Comfort, W.W., Negrepontis, S.: The Theory of Ultrafilters. Springer, Berlin (1974). https://doi.org/10.1007/978-3-642-65780-1