Entanglements

Carmesin, Johannes; Kurkofka, Jan

Journal of Combinatorial Theory
Series B

Notes

Entanglements

Johannes Carmesin [1], Jan Kurkofka [2]

University of Birmingham, Birmingham, UK

A R T I C L E   I N F O

A B S T R A C T

Robertson and Seymour constructed for every graph G a tree-decomposition that efficiently distinguishes all the tangles in G. While all previous constructions of these decompositions are either iterative in nature or not canonical, we give an explicit one-step construction that is canonical.

The key ingredient is an axiomatisation of 'local properties' of tangles. Generalisations to locally finite graphs and matroids are also discussed.

1. Introduction

In this paper we propose an axiomatisation of 'local properties' of tangles and apply it to give explicit one-step constructions of tree-decompositions, as follows.

Roughly speaking, *tree-decompositions* are a recipe how to cut up a graph along separations in a tree-like way. *Tangles* are a way to axiomatise highly cohesive substructures in graphs such as complete subgraphs or grid minors. We say that a separation $\{A, B\}$ of a graph G *distinguishes* a pair of tangles if the two tangles live on opposite sides of $\{A, B\}$; it does so *efficiently* if the separator of $\{A, B\}$ has smallest size amongst all dis-

tinguishing separations of G. We say that a tree-decomposition of a graph G (*efficiently*) *distinguishes* a pair of tangles if there is a separation $\{A, B\}$ which (efficiently) distinguishes the two tangles and $\{A, B\}$ is in the recipe for the tree-decomposition. A key tool [18] in the proof of the graph-minor theorem states:

> *Every finite graph G has a tree-decomposition that efficiently distinguishes*
>
> *all the tangles in G.* (1)

A fair amount of the recent work on graph-minors has focused on constructing such tree-decompositions [1,4–7,9,12,15,17]. In all proofs in the literature these tree-decompositions are constructed through an iterative process in which separations are chosen in turn based on previous choices. Here we will give a new construction of the tree-decomposition of (1) that finishes in one step, is canonical, and that is *explicit* in the sense that it computes a single simple parameter for separations and then takes all separations for the tree-decomposition which minimise this parameter.

In the proof of (1), one has to construct separations that distinguish all pairs of tangles efficiently, and one has to construct them in a *nested* way; that is, so that they define the recipe of a tree-decomposition. Rather than working with tangles in the first place, our perspective is to directly axiomatise separations which distinguish tangles efficiently through a new notion of *entanglements*; see Section 2. Perhaps surprisingly, these entanglements have very similar properties to tangles themselves but only applied to a subset of their separations. See Section 2 for an explanation of why we think of entanglements as an axiomatisation of 'local properties' of tangles.

Our main result reads as follows.

Definition *(Friendly).* A separation $\{A, B\}$ in an entanglement ε in G is *friendly* if no other separation in ε crosses less separations in entanglements in G than $\{A, B\}$.

Theorem 1. *For every finite graph G, the set of friendly separations of G is a nested set of separations; and hence gives rise to a tree-decomposition distinguishing all tangles efficiently.*

The nested sets $N(G)$ and tree-decompositions $\mathcal{T}(G)$ provided by Theorem 1 are canonical in that they commute with graph-isomorphisms: $\varphi(N(G)) = N(\varphi(G))$ and $\varphi(\mathcal{T}(G)) = \mathcal{T}(\varphi(G))$ for every graph-isomorphism $\varphi \colon G \to G'$.

The decomposition in Theorem 1 refines the one of (1). Indeed, not every entanglement is induced by a pair of tangles, and in fact entanglements and friendly separations can be found in graphs that host no tangles at all (Example 2.2).

Theorem 1 extends to locally-finite infinite graphs under additional assumptions; see Theorem 4.2. We also provide an abstract version of Theorem 1, inspired by [10–12,14], which can be applied to a wide variety of setups including matroids; see Section 5.

Fig. 1. $\{A \cap C, B \cup D\}$ is one of the four corners of $\{A, B\}$ and $\{C, D\}$.

This note is organised as follows. Entanglements in graphs are introduced in Section 2. Theorem 1 is proved in Section 3. An infinite version of Theorem 1 is proved in Section 4. Abstract versions of entanglements and of Theorem 1 are offered in Section 5.

2. Entanglements in graphs

Let G be any graph. A *separation* of G is a set $\{A, B\}$ such that $A \cup B = V(G)$ and G contains no edge between $A \setminus B$ and $B \setminus A$. We refer to A and B as the *sides* of $\{A, B\}$, and call $A \cap B$ the *separator* of $\{A, B\}$. The size $|A \cap B|$ of the separator is the *order* of $\{A, B\}$. A separation $\{A, B\}$ is *proper* if $A \setminus B$ and $B \setminus A$ are non-empty. Two separations $\{A, B\}$ and $\{C, D\}$ of G are *nested* if, after possibly renaming their sides, they satisfy $A \subseteq C$ and $B \supseteq D$. Two separations that are not nested are said to *cross*. A set of separations of G is *nested* if its elements are pairwise nested.

For a depiction of the setting for the next definitions, see Fig. 1. If $\{A, B\}$ and $\{C, D\}$ cross, then their four *corners* are the separations $\{A \cap C, B \cup D\}$, $\{A \cap D, B \cup C\}$, $\{B \cap D, A \cup C\}$ and $\{B \cap C, A \cup D\}$. The corners $\{A \cap C, B \cup D\}$ and $\{B \cap D, A \cup C\}$ are *opposite*, and so are the corners $\{A \cap D, B \cup C\}$ and $\{B \cap C, A \cup D\}$. Any two corners that are not opposite are *adjacent*. The two adjacent corners $\{A \cap C, B \cup D\}$ and $\{A \cap D, B \cup C\}$ are said to *lie on the same side* of $\{A, B\}$. Similarly, the two adjacent corners $\{B \cap D, A \cup C\}$ and $\{B \cap C, A \cup D\}$ are said to *lie on the same side* of $\{A, B\}$.

An *entanglement* in G is a non-empty set ε of proper separations of G such that ε satisfies ($\mathcal{E}$):

($\mathcal{E}$) If a separation $\{A, B\} \in \varepsilon$ is crossed by a separation of G so that two corners lying on the same side of $\{A, B\}$ have order at most $|A \cap B|$, then at least one of these corners has order equal to $|A \cap B|$ and is contained in ε.

A separation $\{A, B\}$ in an entanglement ε in G is *friendly* if no other separation in ε crosses less separations in entanglements in G than $\{A, B\}$.

We conclude this section with three examples. The first example uses the terminology of [9, §12.5]. We state it as a lemma because it is a key ingredient of the proof of Theorem 1.

Fig. 2. The Farey graph of order 4. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

Lemma 2.1. *Every pair of distinguishable tangles in a graph induces an entanglement, which consists of the separations efficiently distinguishing the two tangles.*

Proof. Let τ and τ' be two distinguishable tangles in a graph G, and let ε be the set of all separations of G which efficiently distinguish τ and τ'. The set ε is non-empty since τ and τ' are distinguishable, and the separations in ε are proper because tangles do not contain separations of the form $(V(G), B)$. We claim that ε satisfies $(\mathcal{E})$. For this, suppose that $\{A, B\} \in \varepsilon$ is crossed by a separation $\{C, D\}$ of G so that the two corners $c_1 := \{A \cap C, B \cup D\}$ and $c_2 := \{A \cap D, B \cup C\}$ have order at most $|A \cap B|$. Without loss of generality, τ orients $\{A, B\}$ towards A and τ' orients $\{A, B\}$ towards B. Since the corners c_1 and c_2 have order at most $|A \cap B|$, they are oriented by τ and τ'. The tangle τ' orients both c_1 and c_2 towards B by consistency. The tangle τ cannot orient both c_1 and c_2 towards B since tangles do not contain three separations whose small sides together cover G. Therefore, τ orients one of c_1 and c_2 away from B. Then that corner distinguishes τ and τ', and must do so efficiently, hence it lies in ε. $\quad\square$

If τ is a tangle in G, and σ_i for $i \in I$ are the tangles in G that are distinguishable from τ, then for every σ_i we obtain an entanglement $\varepsilon_i \subseteq \tau$ by Lemma 2.1, and these ε_i contain all the information from τ that is sufficient to efficiently distinguish τ from all σ_i. This is why intuitively, we may think of entanglements as an axiomatisation of 'local properties' of tangles.

Example 2.2. The *Farey graph F_1 of order* 1 is obtained from a 4-cycle whose edges are coloured blue by adding a chord. Recursively, the *Farey graph F_{k+1} of order $k + 1$* is obtained from F_k by adding a new vertex v_e for each blue edge e of F_k, joining it to the two endvertices of e with blue edges, and uncolouring the previously blue edge e; see Fig. 2. Now let $k \in \mathbb{N}$ be any number and let us consider F_k.

Each non-blue edge of F_k leaves two components after deleting its endvertices, and therefore defines a separation of F_k in the obvious way. Let N be the set of all separations of F_k defined in this way. We claim that each separation in N forms an entanglement of its own. To see that these singletons satisfy $(\mathcal{E})$, consider any separation $\{A, B\} \in N$, and let $\{C, D\}$ be any separation of G which crosses $\{A, B\}$. It suffices to show that

Fig. 3. The situation in Example 2.3.

of every two corners lying on the same side of $\{A, B\}$, at least one has order larger than $|A \cap B| = 2$. Since $A \cap B$ induces a K^2, we may assume without loss of generality that $A \cap B \subseteq C$. By symmetry, it suffices to show that the corner $c := \{A \cap C, B \cup D\}$ has order at least three. If the separator of c has size at most two, then it is equal to $A \cap B$, and $c = \{A, B\}$ follows because $A \setminus B$ and $B \setminus A$ are connected. In particular, $A \cap C = A$ implies $A \subseteq C$, and $B \cup D = B$ implies $B \supseteq D$, so $\{A, B\}$ and $\{C, D\}$ are nested. Since this would contradict our assumptions, c must have order at least three, as desired. Hence, each separation in N forms an entanglement in F_k, so each separation in N is a friendly separation of F_k.

The set of all separations of F_k whose separators span a K^2 is nested and witnesses that there is no tangle in F_k by [9, Theorem 12.5.1].

Example 2.3. We claim that wheels have no entanglements. Indeed, let G be a wheel and let us suppose for a contradiction that there is an entanglement in G. Let $\{A, B\}$ be a separation of G that lies in an entanglement and whose side A is inclusionwise minimal among all separations of G that lie in entanglements. Since $A \setminus B$ and $B \setminus A$ are non-empty and the centre c of the wheel is joined to all other vertices, c can only be contained in $A \cap B$; see Fig. 3. Furthermore, $|A \cap B| \geqslant 3$ since G is 3-connected. Pick any two vertices $a \in A \setminus B$ and $b \in B \setminus A$, and let $\{C, D\}$ be the separation of G with $C \cap D = \{a, b, c\}$. Let P and Q be the two internally disjoint a–b paths through the rim of the wheel. Since $A \cap B$ meets both P and Q in internal vertices, it follows that $\{A, B\}$ and $\{C, D\}$ cross and that all four corners have order at most $|A \cap B|$. Hence $(\mathcal{E})$ implies that at least one of the corners on the A-side of $\{A, B\}$ lies in an entanglement. This contradicts the minimal choice of A.

3. Friendly separations are nested

For a finite graph G and a separation s of G, let us denote by $x(s)$ the number of separations in entanglements in G which are crossed by s, and call $x(s)$ the *crossing number* of s in G.

Lemma 3.1. *Let G be any finite graph. Suppose that for all entanglements $\varepsilon_1, \varepsilon_2$ in G (possibly with $\varepsilon_1 = \varepsilon_2$) and any two crossing separations $s_1 \in \varepsilon_1$ and $s_2 \in \varepsilon_2$, there exist*

an index $i \in \{1, 2\}$ and a separation $c \in \varepsilon_i$ such that $x(c) < x(s_i)$. Then the friendly separations of G are nested. □

Lemma 3.2. *Let G be any graph, let r, s be two crossing separations of G, and let c, d be two opposite corners of r, s. For every separation t of G the following assertions hold:*

  (i) *If t crosses at least one of c and d, then t crosses at least one of r and s.*
 (ii) *If t crosses both c and d, then t crosses both r and s.*
(iii) *Neither r nor s crosses c or d.*

Proof. (i) holds by [9, Lemma 12.5.5], whose proof works for both finite and infinite graphs. (ii) is straightforward if one shows the contrapositive. (iii) is trivial. □

Corollary 3.3. *Let G be any finite graph, let r, s be two crossing separations in entanglements in G, and let c, d be two opposite corners of r, s. For every separation t of G we have $x(c) + x(d) < x(r) + x(s)$.*

Proof. Combining (i)–(iii) of Lemma 3.2 gives $x(c) + x(d) \leqslant x(r) + x(s)$. Since s and r lie in entanglements and cross, they are counted in $x(r)$ and in $x(s)$, but they contribute to neither $x(c)$ nor $x(d)$ by (iii); hence the inequality is strict. □

Lemma 3.4. *Let G be any finite graph. Suppose that for all entanglements $\varepsilon_1, \varepsilon_2$ in G (possibly with $\varepsilon_1 = \varepsilon_2$) and any two crossing separations $s_1 \in \varepsilon_1$ and $s_2 \in \varepsilon_2$, at least one of the following conditions is satisfied:*

(C1) *there are opposite corners c_1, c_2 of s_1, s_2 with $c_1 \in \varepsilon_1$ and $c_2 \in \varepsilon_2$;*
(C2) *two opposite corners of s_1, s_2 are in ε_1, and the other two opposite corners of s_1, s_2 are in ε_2.*

Then the friendly separations of G are nested.

Proof. It suffices to show that the premise of Lemma 3.1 is satisfied. For this, let $\varepsilon_1, \varepsilon_2$ be any entanglements in G (possibly with $\varepsilon_1 = \varepsilon_2$) and let s_1, s_2 be two crossing separations with $s_1 \in \varepsilon_1$ and $s_2 \in \varepsilon_2$.

First, suppose that by (C1) there are opposite corners c_1, c_2 of s_1, s_2 with $c_1 \in \varepsilon_1$ and $c_2 \in \varepsilon_2$. By Corollary 3.3, we have $x(c_1) + x(c_2) < x(s_1) + x(s_2)$. An indirect proof finds an $i \in \{1, 2\}$ with $x(c_i) < x(s_i)$.

Second, suppose that by (C2) there are two opposite corners c_1, c_1' of s_1, s_2 are in ε_1, and the other two opposite corners c_2, c_2' of s_1, s_2 are in ε_2. By Corollary 3.3, we have $x(c_i) + x(c_i') < x(s_1) + x(s_2)$ for both $i = 1, 2$. Without loss of generality, we have $x(s_1) \leqslant x(s_2)$. Hence either $x(c_2) < x(s_2)$ or $x(c_2') < x(s_2)$. □

Fig. 4. The situation in Case 1.

Let us write $|s| := |A \cap B|$ for a separation $s = \{A, B\}$. If two separations s_1 and s_2 of a graph G cross and c_1, c_2 are two opposite corners of s_1, s_2, then the orders of these corners sum to $|c_1| + |c_2| = |s_1| + |s_2|$. The important part of this equality is the inequality $|c_1| + |c_2| \leqslant |s_1| + |s_2|$, which is known as *submodularity*, and which is the only part of the equality that we will need in the proofs.

Theorem 3.5. *The friendly separations of any finite graph are nested.*

Proof. It suffices to show that the premise of Lemma 3.4 is satisfied. For this, let ε_1 and ε_2 be any entanglements in G, possibly with $\varepsilon_1 = \varepsilon_2$, and let $s_1 \in \varepsilon_1$ and $s_2 \in \varepsilon_2$ be two crossing separations. Without loss of generality, we have $|s_1| \leqslant |s_2|$. Let us colour a corner of s_1, s_2 green if it has order at most $|s_2|$.

Sublemma 3.6. *At least three corners of s_1, s_2 are green.*

Proof. Suppose for a contradiction that at most two corners of s_1, s_2 are green. By submodularity and since $|s_1| \leqslant |s_2|$, at least one of any two opposite corners must be green. So there are exactly two green corners, and since they cannot be opposite they must be adjacent. As the remaining two corners are not green by assumption, they have order greater than $|s_2|$. By submodularity, this means that the green corners in fact have order less than $|s_1|$. Then either ε_1 or ε_2 contains a green corner by $(\mathcal{E})$. But then this green corner has order equal to $|s_1|$ or $|s_2|$ by $(\mathcal{E})$, contradicting our observation that it has order less than $|s_1|$ and $|s_2|$. $\square$

By Sublemma 3.6, at least three corners of s_1, s_2 are green. Hence it suffices to consider the following two cases. See Fig. 4 for a depiction of Case 1.

Case 1: In the first case, precisely three corners of s_1, s_2 are green. Then two green corners c_2, d_2 lie on the same side of s_2, so at least one of them is contained in ε_2 by $(\mathcal{E})$, say $c_2 \in \varepsilon_2$. Hence c_2 has order exactly $|s_2|$. So the corner c_1 opposite of c_2 has order at most $|s_1|$ by submodularity; in particular, c_1 is green. Note that c_1 and d_2 lie on the same side of s_1. The corner opposite of d_2 is not green, so has order more than $|s_2|$. Hence d_2 has order less than $|s_1|$ by submodularity. So by $(\mathcal{E})$, at least one of d_2 and c_1 is contained in ε_1 and has order equal to $|s_1|$. This can only be c_1. So c_1, c_2 are opposite corners of s_1, s_2 with $c_1 \in \varepsilon_1$ and $c_2 \in \varepsilon_2$, giving (C1).

Case 2: In the second case, all four corners are green. Applying $(\mathcal{E})$ on both sides of $s_2 \in \varepsilon_2$, we find corners c_2, c_2' of s_1, s_2 with $c_2, c_2' \in \varepsilon_2$ such that c_2, c_2' do not lie on the same side of s_2. Moreover, c_2 and c_2' have order exactly $|s_2|$ by $(\mathcal{E})$. We consider two subcases.

Subcase 2A: In the first subcase, the two corners c_2, c_2' are adjacent, so they lie on the same side of s_1. Let c_1 be the corner opposite of c_2, and let c_1' be the corner opposite of c_2'. The corners c_1, c_1' have order at most $|s_1|$ by submodularity. Moreover, c_1 and c_1' lie on the same side of s_1. Hence at least one of c_1 and c_1' is contained in ε_1 by $(\mathcal{E})$, and we already know that its opposite corner is contained in ε_2, giving (C1).

Subcase 2B: In the second subcase, the two corners c_2, c_2' are opposite. Since c_2 and c_2' have order $|s_2|$, submodularity with $|s_1| \leqslant |s_2|$ implies $|s_1| = |s_2|$. Therefore, by symmetry we can repeat the entire argumentation up to this point with the roles of s_1 and s_2 interchanged to find two opposite corners c_1, c_1' of s_1, s_2 with $c_1, c_1' \in \varepsilon_1$. If the sets $\{c_1, c_1'\}$ and $\{c_2, c_2'\}$ intersect, then they are equal, so c_2 and c_2' are opposite corners of s_1, s_2 with $c_2 \in \varepsilon_1$ and $c_2' \in \varepsilon_2$, giving (C1). Otherwise, $\{c_1, c_1'\}$ and $\{c_2, c_2'\}$ are disjoint, and then c_1, c_1' are two opposite corners of s_1, s_2 in ε_1 while c_2, c_2' are the other two opposite corners of s_1, s_2 and are in ε_2, giving (C2). □

Proof of Theorem 1. Let G be a finite graph, and let N denote its set of friendly separations. The set N is nested by Theorem 3.5, and it efficiently distinguishes all the tangles in G by Lemma 2.1. As is well-known [18, (9.1)], N defines a tree-decomposition $\mathcal{T}$ of G, which efficiently distinguishes all the tangles in G since N does. □

Remark 3.7. To construct the tree-decomposition $\mathcal{T}$ that efficiently distinguishes all the tangles in the proof of Theorem 1, we have used all entanglements in G (to first define N and then $\mathcal{T}$), not just the ones induced by the pairs of distinguishable tangles. It is possible to adjust the entire framework of this section to only work with the set $\mathcal{E}$ of tangle-induced entanglements instead, to obtain a nested set $N' = N'(\mathcal{E})$, which may be incomparable with N (as set), and then obtain a tree-decomposition from N'; we do this in more detail in Theorem 4.2 (because there we must restrict to a subset of all the entanglements). This would make sure that every separation in N' (and hence of the tree-decomposition) efficiently distinguishes two tangles in G. However, there is an alternative way to achieve this: we can consider the subset $N'' \subseteq N$ formed by the separations that efficiently distinguish some two tangles, and then consider the tree-decomposition defined by N''.

4. Entanglements in locally-finite infinite graphs

Recall that a graph is *locally finite* if each of its vertices has only finitely many neighbours. In this section, we extend Theorem 3.5 to locally-finite infinite graphs. The proof of Theorem 3.5 almost works for locally-finite infinite graphs. The only places where we use finiteness are where we use the crossing numbers $x(s)$; indeed, we only need that

all relevant crossing numbers are finite. To ensure this, we combine local finiteness with two other customary conditions, *tightness* and *finite boundedness*; see Lemma 4.1. Then we extend Theorem 3.5 to infinite graphs under the combination of the three conditions. The combination of the three conditions is mild in the sense that the extension result, Theorem 4.2, is strong enough for its application in [13].

A separation $\{A, B\}$ of a graph G is *tight* if there are components C_A and C_B of $G-(A\cap B)$ with $C_A \subseteq G[A\backslash B]$ and $C_B \subseteq G[B\backslash A]$ such that $N_G(C_A) = A\cap B = N_G(C_B)$. An entanglement in a graph is *tight* if it consists of tight separations. For instance, entanglements induced by pairs of tangles are tight [15, Lemma 6.1].

Lemma 4.1. *Let G be any locally finite connected graph and $k \in \mathbb{N}$. Then every tight finite-order separation of G is crossed by only finitely many tight separations of G of order at most k.*

Proof. This fact is well-known; see e.g. the proof of [15, Proposition 6.2]. $\quad\square$

A set $\mathcal{E}$ of entanglements is *finitely bounded* if there is $k \in \mathbb{N}$ with $|s| \leqslant k$ for all $s \in \bigcup \mathcal{E}$. Let G be any graph, and let $\mathcal{E}$ be a set of entanglements in G. Suppose that $\mathcal{E}$ is finitely bounded. If G is locally finite but infinite, we additionally assume that all entanglements in $\mathcal{E}$ are tight, so that each separation in $\bigcup \mathcal{E}$ crosses only finitely many separations in $\bigcup \mathcal{E}$ by Lemma 4.1. A separation $\{A, B\}$ in an entanglement $\varepsilon \in \mathcal{E}$ is *$\mathcal{E}$-friendly* if no other separation in ε crosses less separations in $\bigcup \mathcal{E}$.

Theorem 4.2. *Let G be any locally-finite connected graph and let $\mathcal{E}$ be any finitely bounded set of tight entanglements in G. Then the set of $\mathcal{E}$-friendly separations of G is nested.*

Proof. The plan is to walk through Section 3 once more and see that everything adjusts to and works in the setting of the theorem. First, we adjust the crossing numbers: $x(s)$ counts only the separations in entanglements in $\mathcal{E}$ that cross s. Then $x(s)$ is finite for all $s \in \bigcup \mathcal{E}$, by Lemma 4.1.

In Lemma 3.1, we only consider entanglements in $\mathcal{E}$, and use that the crossing-numbers $x(s_i)$ are finite. Lemma 3.2 is stated and proved for arbitrary graphs. In Corollary 3.3, we only consider entanglements in $\mathcal{E}$, so $x(r)$ and $x(s)$ are finite; then the proof extends. Lemma 3.4 extends similarly, and so does the proof of Theorem 3.5. $\quad\square$

Recall that every end of a graph induces a tangle of infinite order; in particular, every pair of ends induces an entanglement. Two ends of a graph are *$(<k)$-distinguishable* (for $k \in \mathbb{N}$) if their induced tangles are distinguished by a separation of order less than k.

Corollary 4.3. *Let G be any locally-finite connected graph and $k \in \mathbb{N}$. Let $\mathcal{E}$ be the set of all entanglements in G that are induced by pairs of $(<k)$-distinguishable ends of G. Then the set of $\mathcal{E}$-friendly separations of G is nested and efficiently distinguishes every pair of $(<k)$-distinguishable ends of G.* $\quad\square$

Rühmann showed a result that is somewhat similar to the above corollary, see [19, Theorem 6.1.6]. For more on infinite trees of tangles, we refer to [8,2,15–17].

5. Abstract entanglements

In this section, we introduce an abstract setting which is more general than separations of graphs, and generalise Theorem 1 to this abstract setting.

A *separation* is a set of the form $\{A, B\}$ with $A \neq B$. We refer to A and B as the *(opposite) sides* of $\{A, B\}$. An *uncrossing-setting* on a set S of separations is a pair $(S, \sim)$ where $\sim$ is an anti-reflexive symmetric binary-relation on S. Instead of writing $r \sim s$ we say that r and s *cross*, and any two elements of S that do not cross are *nested*. A set of separations in S is *nested* if its elements are pairwise nested.

A *corner-map* for an uncrossing-setting $(S, \sim)$ is a map $\boxplus$ which assigns to every unordered pair of crossing separations $r = \{A, B\}$ and $s = \{C, D\}$ four pairwise distinct separations $L_{\{r,s\}}(\{X, Y\})$, one for each choice of sides $X \in \{A, B\}$ and $Y \in \{C, D\}$, subject to condition (F) below. We allow any number of these corners to be elements of S, but we do not require them to be elements of S. A corner of r, s that is contained in S shall be called an *S-corner* for emphasis. As r and s will always be clear from context, we reduce the notation $L_{\{r,s\}}(\{X, Y\})$ to $L(X, Y)$ for convenience.

Example 5.1. If two separations $\{A, B\}$ and $\{C, D\}$ of a graph G cross, then the four corners are the usual corners $L(X, Y) := \{X \cap Y, X' \cup Y'\}$ for $\{X, X'\} = \{A, B\}$ and $\{Y, Y'\} = \{C, D\}$.

Two distinct corners $L(X, Y)$ and $L(X', Y')$ are *opposite* if X, X' are opposite sides of $\{A, B\}$ and Y, Y' are opposite sides of $\{C, D\}$. They are *adjacent* if they are not opposite, which is equivalent to having $X = X'$ or $Y = Y'$ but not both. They *lie on the same side* of $\{A, B\}$ if $X = X'$, and similarly they *lie on the same side* of $\{C, D\}$ if $Y = Y'$. Note that distinct corners that lie on the same side of r or of s are adjacent. Condition (F) generalises Corollary 3.3 and reads as follows:

(F) Every two opposite S-corners c, d of r, s satisfy the following three conditions.

   (F1) If $t \in S$ crosses at least one of c and d, then t crosses at least one of r and s.
   (F2) If $t \in S$ crosses both c and d, then t crosses both r and s.
   (F3) Neither r nor s crosses c or d.

An *order-function* is a map

$$| \cdot | \colon S \cup \{\text{corners of crossing separations in } S\} \to \mathbb{R}_{\geqslant 0}.$$

Then $|s|$ is the *order* of s. An order-function $| \cdot |$ is *submodular* if for every two crossing elements $r, s \in S$ and opposite corners c, d of r, s it satisfies $|c| + |d| \leqslant |r| + |s|$. A *sub-*

modular uncrossing-setting on a set S of separations is a triple $(S, \sim, \boxplus, |\cdot|)$ formed by an uncrossing-setting $(S, \sim)$ with a corner-map $\boxplus$ and a submodular order-function $|\cdot|$.

An *entanglement* in a submodular uncrossing-setting on a set S of separations is a non-empty subset $\varepsilon \subseteq S$ which exhibits the following property:

($\mathcal{E}$) If a separation $r \in \varepsilon$ is crossed by an $s \in S$ so that two adjacent corners on the same side of r have order at most $|r|$, then at least one of these two corners has order equal to $|r|$ and lies in ε.

Suppose now that S is finite. For every $s \in S$ we denote by $x(s)$ the number of separations in entanglements which are crossed by s, and we call $x(s)$ the *crossing-number* of s.

Lemma 5.2. *Let r, s be two crossing separations in entanglements in a submodular uncrossing-setting on a set S of separations. Then for every two opposite S-corners c, d of r, s we have $x(c) + x(d) < x(r) + x(s)$.*

Proof. This follows from (F1)–(F3) just like in the proof of Corollary 3.3.   □

A separation $s \in S$ is *friendly* if it occurs in an entanglement ε and no other separation in ε crosses less separations in entanglements.

Theorem 5.3. *The friendly separations in a finite submodular uncrossing-setting are nested.*

Proof. The proof is analogous to the proof of Theorem 3.5, including Lemma 3.1 and Lemma 3.4, with just one exception: instead of Corollary 3.3, we use Lemma 5.2.   □

Theorem 5.3 clearly implies Theorem 1, and it yields the following version of Theorem 1 for matroids. We state the theorem using the terminology of [12, §4.2]. The usual order-function for matroid-separations is well known to be submodular, see e.g. [12]. Matroid-separations exhibit (F): indeed, the proof of Lemma 3.2 extends to matroid-separations verbatim. Hence matroid-separations form a submodular uncrossing-setting.

Theorem 5.4. *For every finite matroid M, the set of friendly separations of M is a nested set of separations; and hence gives rise to a tree-decomposition distinguishing all tangles efficiently.*   □

Concluding remarks   In [3], r-local 2-separations of graphs have been introduced, which need not separate the graph globally but which separate it r-locally in that they separate a ball of radius $r/2$ around their separators. While it is not obvious how the notion of tangles could be generalised to r-local separations, this can be achieved for entanglements with a slightly different notion of r-local separations, as announced in [3]. We would also

like to mention that Theorem 4.2 and Theorem 5.3 will be used in upcoming work to find graph-decompositions, see for example [13].

Data availability

No data was used for the research described in the article.

Acknowledgment

References

[1] J. Carmesin, A short proof that every finite graph has a tree-decomposition displaying its tangles, Eur. J. Comb. 58 (2016).

[2] J. Carmesin, All graphs have tree-decompositions displaying their topological ends, Combinatorica 39 (2019) 545–596.

[3] J. Carmesin, Local 2-separators, J. Comb. Theory, Ser. B 156 (2022) 101–144.

[4] J. Carmesin, R. Diestel, M. Hamann, F. Hundertmark, Canonical tree-decompositions of finite graphs I. Existence and algorithms, J. Comb. Theory, Ser. B 116 (2016) 1–24.

[5] J. Carmesin, R. Diestel, M. Hamann, F. Hundertmark, Canonical tree-decompositions of finite graphs II. Essential parts, J. Comb. Theory, Ser. B 118 (2016) 268–283.

[6] J. Carmesin, R. Diestel, F. Hundertmark, M. Stein, Connectivity and tree structure in finite graphs, Combinatorica 34 (1) (2014) 1–35.

[7] J. Carmesin, P. Gollin, Canonical tree-decompositions of a graph that display its k-blocks, J. Comb. Theory, Ser. B 122 (2017) 1–20.

[8] J. Carmesin, M. Hamann, B. Miraftab, Canonical trees of tree-decompositions, J. Comb. Theory, Ser. B 152 (2022) 1–26.

[9] R. Diestel, Graph Theory, 5th edition, Springer, 2017.

[10] R. Diestel, Abstract separation systems, Order 35 (2018) 157–170.

[11] R. Diestel, Tree sets, Order 35 (2018) 171–192.

[12] R. Diestel, F. Hundertmark, S. Lemanczyk, Profiles of separations: in graphs, matroids, and beyond, Combinatorica 39 (1) (2019) 37–75.

[13] R. Diestel, R.W. Jacobs, P. Knappe, J. Kurkofka, Canonical graph decompositions via coverings, arXiv:2207.04855, 2022, submitted for publication.

[14] C. Elbracht, J. Kneip, M. Teegen, Trees of tangles in abstract separation systems, J. Comb. Theory, Ser. A 180 (2021) 105425.

[15] C. Elbracht, J. Kneip, M. Teegen, Trees of tangles in infinite separation systems, Math. Proc. Camb. Philos. Soc. (2021) 1–31, arXiv:1909.09030.

[16] A.K. Elm, J. Kurkofka, A tree-of-tangles theorem for infinite tangles, Abh. Math. Semin. Univ. Hamb. 92 (2022) 139–178.

[17] R.W. Jacobs, P. Knappe, Efficiently distinguishing all tangles in locally finite graphs, arXiv:2303.09332, 2023, submitted for publication.

[18] N. Robertson, P.D. Seymour, Graph minors. X. Obstructions to tree-decomposition, J. Comb. Theory, Ser. B 52 (1991) 153–190.

[19] T. Rühmann, A study of infinite graphs of a certain symmetry and their ends, PhD thesis, Universität Hamburg, 2017.