

Canonical tree-decompositions of finite graphs

I. Existence and algorithms

J. Carmesin R. Diestel M. Hamann F. Hundertmark

April 24, 2014

Abstract

We construct tree-decompositions of graphs that distinguish all their k -blocks and tangles of order k , for any fixed integer k . We describe a family of algorithms to construct such decompositions, seeking to maximize their diversity subject to the requirement that they commute with graph isomorphisms. In particular, all the decompositions constructed are invariant under the automorphisms of the graph.

1 Introduction

A k -block in a graph G , where k is any positive integer, is a maximal set X of at least k vertices such that no two vertices $x, x' \in X$ can be separated in G by fewer than k vertices other than x and x' . Thus, k -blocks can be thought of as highly connected pieces of a graph, but their connectivity is measured not in the subgraph they induce but in the ambient graph.

Extending results of Tutte [9] and of Dunwoody and Krön [5], three of us and Maya Stein showed that every finite graph G admits, for every integer k , a tree-decomposition (T, \mathcal{V}) of adhesion $< k$ that distinguishes all its k -blocks [3]. These decompositions are *canonical* in that the map $G \mapsto (T, \mathcal{V})$ commutes with graph isomorphisms. In particular, the decomposition (T, \mathcal{V}) constructed for G is invariant under the automorphisms of G .

Our next aim, then, was to find out more about the tree-decompositions whose existence we had just proved. What can we say about their parts? Will every part contain a k -block? Will those that do consist of just their k -block, or might they also contain some ‘junk’? Such questions are not only natural; their answers will also have an impact on the extent to which our tree-decompositions can be used for an obvious potential application, to the graph isomorphism problem in complexity theory. See Grohe and Marx [6] for recent progress on this.

When we analysed our existence proof in view of these questions, we found that even within the strict limitations imposed by canonicity we can make choices that will have an impact on the answers. For example, we can obtain different decompositions (all canonical) if we seek to, alternatively, minimize the number of inessential parts, minimize the sizes of the parts, or just of the essen-

tial parts, or achieve a reasonable balance between these properties. (A part is called *essential* if it contains a k -block, and *inessential* otherwise.)

In this paper we describe a large family of algorithms¹ that each produce a canonical tree-decomposition for given G and k . Their parameters can be tuned to optimize this tree-decomposition in terms of criteria such as those above. In [1] we shall apply these results to specify algorithms from the family described here for which we can give sharp bounds on the number of inessential parts, or which under specified conditions ensure that some or all essential parts consist only of the corresponding k -block.

The existence theorems which our algorithms imply will extend our results from [3] in that the decompositions constructed will not only distinguish all the k -blocks of a graph, but also its tangles of order k . (Tangles were introduced by Robertson and Seymour [8] and can also be thought of as indicating highly connected parts of a graph.) In order to treat blocks and tangles in a unified way, we work with a common generalization called ‘profiles’. These appear to be of interest in their own right, as a way of locating desirable local substructures in very general discrete structures. More about profiles, including further generalizations of our existence theorems to such general structures (including matroids), can be found in [7]. More on k -blocks, including different kinds of examples, their relationship to tangles, the algorithmic complexity of finding them, and a block-decomposition duality theorem, can be found in [2].

All graphs in this paper will be finite, undirected and simple. Any graph-theoretic terms not defined here are explained in [4]. Unless otherwise mentioned, G will denote an arbitrary finite graph with vertex set V .

2 Separation systems

A pair (A, B) of subsets of V such that $A \cup B = V$ is called a *separation* of G if there is no edge $e = \{x, y\}$ in G with $x \in A \setminus B$ and $y \in B \setminus A$. A separation (A, B) is *proper* if neither $A \subseteq B$ nor $B \subseteq A$; otherwise it is *improper*. The *order* of a separation (A, B) is the cardinality of its *separator* $A \cap B$. A separation of order k is a *k -separation*. By a simple calculation we obtain:

Lemma 2.1. *For any two separations (A, B) and (C, D) , the orders of the separations $(A \cap C, B \cup D)$ and $(B \cap D, A \cup C)$ sum to $|A \cap B| + |C \cap D|$. \square*

We define a partial ordering on the set of separations of G by

$$(A, B) \leq (C, D) :\Leftrightarrow A \subseteq C \wedge B \supseteq D. \quad (1)$$

¹We should point out that our reason for thinking in terms of algorithms is not, at this stage, one of complexity considerations: these are interesting, but they are not our focus here. Describing a decomposition in terms of the algorithm that produces it is simply the most intuitive way to ensure that it will be canonical: as long as the instructions of how to obtain the decomposition refer only to invariants of the graph (rather than, say, to a vertex enumeration that has to be chosen arbitrarily at some point), the decomposition that this algorithm produces will also be an invariant.

A separation (A, B) is *nested* with a separation (C, D) , written as $(A, B) \parallel (C, D)$, if it is \leq -comparable with either (C, D) or (D, C) . Since

$$(A, B) \leq (C, D) \Leftrightarrow (D, C) \leq (B, A), \quad (2)$$

the relation \parallel is reflexive and symmetric.² Two separations that are not nested are said to *cross*.

A separation (A, B) is *nested* with a set S of separations, written as $(A, B) \parallel S$, if $(A, B) \parallel (C, D)$ for every $(C, D) \in S$. A set S of separations is *nested* with set S' of separations, written as $S \parallel S'$, if $(A, B) \parallel S'$ for every $(A, B) \in S$; then also $(C, D) \parallel S$ for every $(C, D) \in S'$.

A set of separations is called *nested* if every two of its elements are nested. It is called *symmetric* if whenever it contains a separation (A, B) it also contains (B, A) , and *antisymmetric* if it contains no separation (A, B) together with its *inverse* (B, A) . A symmetric set of separations is a *system* of separations, or *separation system*; if all its separations are proper, it is a *proper separation system*.

A separation (A, B) *separates* a set $X \subseteq V$ if X meets both $A \setminus B$ and $B \setminus A$. Given a set S of separations, we say that X is *S -inseparable* if no separation in S separates X . An *S -block* of G is a maximal S -inseparable set of vertices.

Recall that a *tree-decomposition* of G is a pair (T, \mathcal{V}) of a tree T and a family $\mathcal{V} = (V_t)_{t \in T}$ of vertex sets $V_t \subseteq V(G)$, one for every node of T , such that:

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

(T2) for every edge $e \in G$ there exists a $t \in T$ such that both ends of e lie in V_t ;

(T3) $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever t_2 lies on the t_1 - t_3 path in T .

The sets V_t in such a tree-decomposition are its *parts*. Their intersections $V_t \cap V_{t'}$ for edges tt' of the *decomposition tree* T are the *adhesion sets* of (T, \mathcal{V}) ; their maximum size is the *adhesion* of (T, \mathcal{V}) .

Deleting an oriented edge $\vec{e} = t_1 t_2$ of T divides $T - e$ into two components $T_1 \ni t_1$ and $T_2 \ni t_2$. Then $(\bigcup_{t \in T_1} V_t, \bigcup_{t \in T_2} V_t)$ is a separation of G with separator $V_{t_1} \cap V_{t_2}$ [4, Lemma 12.3.1]; we say that our edge \vec{e} *induces* this separation. A node $t \in T$ is a *hub node* if the corresponding part V_t is the separator of a separation induced by an edge of T at t . If t is a hub node, we call V_t a *hub*.

As is easy to check, the separations induced by (the edges of T in) a tree-decomposition (T, \mathcal{V}) are nested. Conversely, we proved in [3] that every nested separation system³ is induced by some tree-decomposition:

Theorem 2.2. [3, Theorem 4.8] *Every nested proper separation system N is induced by a tree-decomposition (T, \mathcal{V}) of G such that*

- (i) *every N -block of G is a part of the decomposition;*
- (ii) *every part of the decomposition is either an N -block of G or a hub.*

See [3] for how these tree-decompositions are constructed from N .

²But it is not in general transitive, compare [3, Lemma 2.2].

³In [3], as here, we needed this only for proper separation systems. However, the result and proof remain valid for arbitrary nested separation systems.

Let k be a positive integer. A set I of at least k vertices is called ($< k$)-*inseparable* if it is S -inseparable for the set $S = S_k$ of all separations of order $< k$, that is, if for every separation $(A, B) \in S_k$ we have either $I \subseteq A$ or $I \subseteq B$. A maximal ($< k$)-inseparable set of vertices is a k -*block* of G .

Since a k -block is too large to be contained in the separator $A \cap B$ of a separation (A, B) of order $< k$, it thus ‘chooses’ one of the sides A or B , the one containing it. Equivalently, every k -block b chooses one of each inverse pair (A, B) and (B, A) of separations in S_k if we want it to lie in the right-hand side of that separation. Let us give the set of separations chosen by b a name:

$$P_k(b) := \{ (A, B) : |A \cap B| < k \wedge b \subseteq B \}. \quad (3)$$

Another way of making informed choices for small-order separations, but one that cannot necessarily be defined by setting a target set like b , are *tangles*. Tangles were introduced by Robertson and Seymour [8]; the definition can also be found in [4]. Like k -blocks, k -tangles (those of order k) have been considered as a way of identifying the ‘ k -connected components’ of a graph.

In order to treat blocks and tangles together in a unified way, let us distill the common essence of the ‘informed choices’ they define for small-order separations into a couple of axioms, and then just work with these.

One common property of k -tangles P and the sets $P = P_k(b)$ defined by a k -block is

$$(A, B) \in P \wedge (C, D) \leq (A, B) \Rightarrow (D, C) \notin P.$$

We shall call sets P of separations that satisfy this implication *consistent*.⁴ Note that consistent sets of separations are antisymmetric: since $(A, B) \leq (A, B)$, a consistent set P cannot contain both (A, B) and its inverse (B, A) .

Another common property of k -tangles P and the sets $P = P_k(b)$ is this:

$$\text{For all } (A, B), (C, D) \in P \text{ we have } (B \cap D, A \cup C) \notin P. \quad (P)$$

If we think of the right-hand side B of (A, B) , the side to which it ‘points’, as its large side, condition (P) becomes reminiscent of the property of (ultra)filters that ‘the intersection of large sets are large’. An important difference is that rather than demanding that $(A \cup C, B \cap D) \in P$, condition (P) only asks that the converse of this separation shall not be in P . Similarly, consistent sets of separations have a property reminiscent of an ultrafilter being non-principal:

$$\text{If } P \text{ is consistent, it contains no separation of the form } (V, A). \quad (4)$$

Indeed, as $(A, V) \leq (V, A)$, a consistent set of separations containing (V, A) must not contain the inverse of (A, V) , which is (V, A) .⁵ Note that consistent sets of separations can contain improper separations of the form (A, V) .

⁴Note that we do not require that (C, D) must lie in P , only that (D, C) does not. The term ‘consistent’ is natural if we think of a separation (A, B) as ‘pointing towards’ B , so that all the elements of $P_k(b)$ point towards b : then (D, C) points away from D , which makes it inconsistent with (A, B) , which points towards B , if $(C, D) \leq (A, B)$ and hence $D \supseteq B$.

⁵If P is an orientation of S_k (see below) and satisfies (P), the condition that $(V, A) \notin P$ for all $A \subseteq V$ is in fact equivalent to the consistency of P .

Note that while consistency says something only about nested separations, condition (P) is essentially about crossing separations. A consistent set of separations satisfying (P) will be called a *profile*. Consistent sets of separations that do not satisfy (P) will play a role later, too.⁶

Since consistency and condition (P) only require the absence of certain separations from P , a requirement easy to meet by making P small, profiles, unlike blocks and tangles, do not as such witness the existence of any highly connected substructure in a graph.⁷ But they do as soon as we make them large, e.g. by requiring that they contain one of every pair of inverse separations in S_k . To give such rich profiles a name, let us call a profile P a *k-profile* if it satisfies

$$\begin{aligned} & \text{Every separation in } P \text{ has order } < k, \text{ and for every separation} \\ & (A, B) \text{ of order } < k \text{ exactly one of } (A, B) \text{ and } (B, A) \text{ lies in } P. \end{aligned} \quad (5)$$

By (4), every k -profile contains every separation (A, V) with $|A| < k$. As we have seen, the set $P_k(b)$ in (3) is a k -profile; we call it the k -profile *induced by*, or simply *of*, b . A k -profile induced by some k -block is a *k-block profile*.

Since a k -block is a maximal ($< k$)-inseparable set of vertices, there is for every pair of distinct k -blocks b, b' a separation (A, B) of order $< k$ such that $(A, B) \in P_k(b)$ and $(B, A) \in P_k(b')$ [3, Lemma 2.1]. Hence $P_k(b) \neq P_k(b')$. Thus, while every k -block induces a k -profile, conversely a k -profile P is induced by at most one k -block, which we then denote by $b(P)$. All k -block profiles P then satisfy $P = P_k(b(P))$, and we say that b and P *correspond*.

Not every k -profile is induced by a k -block. For example, there are tangles of order k that are not induced by a k -block, such as the unique tangle of any order $k \geq 5$ in a large grid (which has no k -block for $k \geq 5$; see [2, Example 3]). Conversely, there are k -block profiles that are not tangles; indeed, there are graphs that have interesting k -block profiles but have no non-trivial tangle at all [2, Examples 4–5 and Section 6]. The notion of a k -profile thus unifies the ways in which k -blocks and tangles of order k choose one side of every separation of order $< k$, but neither of these two instances of k -profiles generalizes the other.

Let S be any set of separations of G . An S -block X of G is called *large* (with respect to S) if it is not contained in the separator of a separation in S . If all the separations in S have order $< k$, an obvious but typical reason for an S -block to be large is that it has k or more vertices. In analogy to (3) we define for a large S -block X

$$P_S(X) := \{(A, B) \in S \mid X \subseteq B\} \subseteq S. \quad (6)$$

⁶As a typical example, consider the union of three large complete graphs X_1, X_2, X_3 identified in a common triangle. The three 3-separations whose left side is one of X_1, X_2, X_3 are consistent but do not satisfy (P), because the separation $(B \cap D, A \cup C)$ in (P) happens to be one of the original three separations. The analogous system with four complete graphs does satisfy (P).

⁷Readers familiar with the notion of *preferences*, or *havens* – a way of making consistent choices of components of $G - X$ for small sets X of vertices – will recognize this: it is because a preference or haven assigns a component of $G - X$ to *every* set X of $< k$ vertices for some k that the bramble formed by these components has order $\geq k$.

Clearly, $P_S(X)$ is a profile; we call it the *S-profile of X*. As before, the *S*-profiles $P_S(X)$ and $P_S(X')$ of distinct large *S*-blocks X, X' are distinct.

Not every *k*-profile has this form. For example, a tangle θ of order $k \geq 5$ in a large grid is not the *S*-profile of a large *S*-block X for any set $S \supseteq \theta$ of separations, since X would be contained in a large θ -block but the grid has none.

Although profiles are, formally, sets of separations, our intuition behind them is that they signify some ‘highly connected pieces’ of our graph G . Our aim will be to separate all these pieces in a tree-like way, and we shall therefore have to speak about sets of separations that, initially, are quite distinct from the profiles they are supposed to separate. To help readers keep their heads in this unavoidable confusion, we suggest that they think of the sets S of separations discussed below as (initially) quite independent of the profiles P discussed along with them, the aim being to explore the relationship between the two.

A separation (A, B) *distinguishes* two subsets of V if one lies in A , the other in B , and neither in $A \cap B$. A set S of separations *distinguishes* two sets of vertices if some separation in S does.

A separation (A, B) *distinguishes* two sets P, P' of separations if $(A, B) \in P \setminus P'$ and $(B, A) \in P' \setminus P$, or vice versa. Thus, a $(< k)$ -separation distinguishes two *k*-blocks if and only if it distinguishes their *k*-profiles. A set of separations S *distinguishes* P from P' if some separation in S distinguishes them, and S *distinguishes* a set \mathcal{P} of sets of separations if it distinguishes every two elements of \mathcal{P} . By (4),

Only proper separations can distinguish two consistent sets of separations, e.g., two profiles. (7)

An antisymmetric set P of separations *orients* a set S of separations if, for every $(A, B) \in S$, either $(A, B) \in P$ or $(B, A) \in P \cap S$. We then call $P \cap S$ an *orientation* of S , and an *S-profile* if it is a profile.

Distinct consistent orientations O, O' of S are distinguished by some separation in S . Indeed, as O and O' are distinct we may assume that there is a separation $(A, B) \in O \setminus O'$. Then $(B, A) \notin O$, since O is consistent, and $(B, A) \in O'$, since $(A, B) \in S$ and O' orients S . Hence (A, B) distinguishes O from O' .

If P orients S and some set X of vertices lies in B for every $(A, B) \in P \cap S$, we say that P *orients S towards X*. If P is a profile then so is $P \cap S$; we call it the *S-profile of P*.⁸

A profile orienting a set S of separations need not orient it towards any non-empty set of vertices: consider, for example, our earlier tangle θ with $S = \theta$. However, a consistent set P orienting a *nested* set N of separations always orients it towards the intersection X of all the sides B of separations $(A, B) \in P \cap N$. By consistency, this set contains all the separators of the \leq -maximal separations in $P \cap N$, so it is non-empty if G is connected. If X is an *N*-block, we say that P *lives in this N-block X*.

⁸This formalizes the idea that P , thought of as a big chunk of G , lies on exactly one side of every separation in S . For example, if the separations in S have order $< k$ and S is symmetric, then every *k*-profile will orient S .

However, X need not be an N -block: although it is N -inseparable, it can be contained in a separator of a separation in N and extend to more than one N -block of G ; then P will not orient N towards any of these. This can happen even if P is a profile,⁹ but not if P is a k -profile: then X cannot be contained in the separator of any separation in $P \cap N$. This is interesting, because it makes k -profiles of nested separation systems ‘big’ in a way arbitrary profiles need not be. So if P is a k -profile, then X is an N -block towards which P orients N .¹⁰

Given a set S of separations of G and a set \mathcal{P} of profiles orienting S , let us say that two profiles $P, P' \in \mathcal{P}$ *agree on S* if their S -profiles coincide, that is, if $P \cap S = P' \cap S$. This is an equivalence relation on \mathcal{P} , whose classes we call the S -blocks of \mathcal{P} . By definition, elements P, P' of the same S -block \mathcal{Q} of \mathcal{P} have the same S -profile $P \cap S = P' \cap S$, which we call the S -profile of \mathcal{Q} .

A separation (A, B) *splits* a consistent set P of separations if both $P \cup \{(A, B)\}$ and $P \cup \{(B, A)\}$ are consistent. (This implies that neither (A, B) nor (B, A) is in P .) For example, the S -profile corresponding to an S -block \mathcal{Q} of a set \mathcal{P} of profiles orienting a separation system S is split by every separation (A, B) that distinguishes some distinct profiles in \mathcal{Q} . By (4), every separation splitting a consistent set of separations must be proper.

We shall need the following lemma.

Lemma 2.3. *Let N be a nested separation system.*

- (i) *Every proper separation $(A, B) \notin N$ that is nested with N splits a unique consistent orientation O of N . This set O is given by*

$$O = \{(C, D) \in N \mid (C, D) \leq (A, B)\} \cup \{(C, D) \in N \mid (C, D) \leq (B, A)\}.$$

- (ii) *If two separations not contained in but nested with N split distinct consistent orientations of N , they are nested with each other.*

Proof. (i) Since (A, B) is nested with N , for every separation $(C, D) \in N$ either (C, D) or (D, C) is smaller than one of (A, B) or (B, A) and thus contained in

$$O := \{(C, D) \in N \mid (C, D) \leq (A, B)\} \cup \{(C, D) \in N \mid (C, D) \leq (B, A)\}.$$

By definition, O contains only separations from N . As we have seen, every separation from N or its inverse lies in O . Once we know that O is consistent it will follow that O is antisymmetric, so it will be an orientation of N .

To check that O is consistent, consider separations $(E, F) \leq (C, D)$ with $(C, D) \in O$. Our aim is to show that $(F, E) \notin O$. This is clearly the case if $(E, F) \notin N$, since $O \subseteq N$ and N is symmetric, so we assume that $(E, F) \in N$. By definition of O , either $(C, D) \leq (A, B)$ or $(C, D) \leq (B, A)$; we assume the former. Then by transitivity $(E, F) \leq (A, B)$, and hence $(E, F) \in O$ by definition of O . To show that $(F, E) \notin O$ we need to check that $(F, E) \not\leq (A, B)$

⁹For example, let $Z \subseteq V$ be such that $G - Z$ has four components C_1, \dots, C_4 . Let $A_i = V(C_i) \cup Z$ and $B_i = V \setminus V(C_i)$. Then $P = \{(A_i, B_i) \mid i = 1, \dots, 4\}$ is a nested profile, with $X = Z = A_i \cap B_i$ for all i .

¹⁰Conversely, if P orients N towards any N -block, then this is clearly equal to X .

and $(F, E) \not\leq (B, A)$. If $(F, E) \leq (A, B)$ then $(B, A) \leq (E, F) \leq (A, B)$ and hence $B \subseteq A$, contradicting our assumption that (A, B) is proper. If $(F, E) \leq (B, A)$ then $(A, B) \leq (E, F) \leq (A, B)$ and hence $(E, F) = (A, B)$, contradicting our assumption that $(A, B) \notin N$.

So O is a consistent orientation of N , and in particular antisymmetric. By the definition of O , this implies that $O \cup \{(A, B)\}$ and $O \cup \{(B, A)\}$ are consistent. Hence (A, B) splits O , as desired.

It remains to show that O is unique. Suppose (A, B) also splits a consistent orientation $O' \neq O$ of N . Let $(C, D) \in N$ distinguish O from O' , with $(C, D) \in O$ and $(D, C) \in O'$ say. By definition of O , either $(C, D) \leq (A, B)$ or $(C, D) \leq (B, A)$. In the first case $O' \cup \{(A, B)\}$ is inconsistent, since $(B, A) \leq (D, C) \in O' \cup \{(A, B)\}$ but also $(A, B) \in O' \cup \{(A, B)\}$. In the second case, $O' \cup \{(B, A)\}$ is inconsistent, since $(A, B) \leq (D, C) \in O' \cup \{(B, A)\}$ but also $(B, A) \in O' \cup \{(B, A)\}$.

(ii) Consider separations $(A, B), (A', B') \notin N$ that are both nested with N . Assume that (A, B) splits a consistent orientation O of N , and that (A', B') splits a consistent orientation $O' \neq O$ of N . From (4) we know that (A, B) and (A', B') must be proper separations, so they satisfy the premise of (i) with respect to O and O' . As $O \neq O'$, there is a separation $(C, D) \in N$ with $(C, D) \in O$ and $(D, C) \in O'$. By the descriptions of O and O' in (i), the separation (C, D) is smaller than (A, B) or (B, A) , and (D, C) is smaller than (A', B') or (B', A') . The latter is equivalent to (C, D) being greater than (B', A') or (A', B') . Thus, (B', A') or (A', B') is smaller than (C, D) and hence than (A, B) or (B, A) , so (A, B) and (A', B') are nested. \square

We remark that the consistent set O in Lemma 2.3 (i) is usually an N -profile; it is not hard to construct pathological cases in which O fails to satisfy (P), but such cases are rare.

Finally, let us note that every separation system S has a consistent orientation. Indeed, the consistent orientations of S are precisely the orientations P of S that are closed down under \leq , those that contain $(C, D) \in S$ whenever $(C, D) \leq (A, B)$ with $(A, B) \in P$. And we can obtain such an orientation of S greedily: consider any pair of proper separations $\{(A, B), (B, A)\} \subseteq S \setminus P$, and include in P both (A, B) (say) and every separation $(C, D) \leq (A, B)$ from S . The resulting set P is clearly closed down in S under \leq , and remains so if we add all improper separations of the form $(A, V) \in S$. It is easy to check that our collection of separations is now an orientation of S , in particular, that we never included the inverse of a separation added to P previously.

3 Tasks and strategies

In this section we describe a systematic approach to distinguishing some or all of the k -profiles in G by (the separations induced by) canonical tree-decompositions of adhesion less than k . Since the separations induced by a tree-decomposition are nested, our main task in finding such a tree-decomposition will be to

select from the set S of all proper ($< k$)-separations of G , which distinguishes the set of all k -profiles by (4), a nested subset N that will still distinguish all the k -profiles under consideration.

We begin by formalizing the notion of such ‘tasks’. We then show how to solve ‘feasible’ tasks in various ways, and give examples showing how different strategies – all canonical in that they commute with graph isomorphisms – can produce quite different solutions.

Consider a proper separation system S and a set \mathcal{P} of profiles in G . Let us call the pair (S, \mathcal{P}) a *task* if every profile in \mathcal{P} orients S and S distinguishes \mathcal{P} . Another task (S', \mathcal{P}') is a *subtask* of the task (S, \mathcal{P}) if $S' \subseteq S$ and $\mathcal{P}' \subseteq \mathcal{P}$.

The two conditions in the definition of a task are obvious minimum requirements which S and \mathcal{P} must satisfy before it makes sense to look for a nested subset $N \subseteq S$ that distinguishes \mathcal{P} . But to ensure that N exists, S must also be rich enough (in terms of \mathcal{P}): the more profiles we wish to separate in a nested way, the more separations will we need to have available. For example, if S consists of two crossing separations (A, B) , (C, D) and their inverses, and \mathcal{P} contains the four possible orientations of S (which are clearly profiles), then S distinguishes \mathcal{P} but is not nested, while the two subsystems $\{(A, B), (B, A)\}$ and $\{(C, D), (D, C)\}$ of S are nested but no longer distinguish \mathcal{P} . But if we enrich S by adding two ‘corner separations’ $(A \cap C, B \cup D)$, $(A \cup C, B \cap D)$ and their inverses, then these together with (A, B) and (B, A) (say) form a nested subsystem that does distinguish \mathcal{P} .

More generally, we shall prove in this section that we shall be able to find the desired N if S and \mathcal{P} satisfy the following condition:

$$\text{Whenever } (A, B), (C, D) \in S \text{ cross and there exist } P, P' \in \mathcal{P} \text{ such} \\ \text{that } (A, B), (C, D) \in P \text{ and } (B, A), (D, C) \in P', \text{ there exists a sep-} \quad (8) \\ \text{aration } (E, F) \in P \cap S \text{ such that } (A \cup C, B \cap D) \leq (E, F).$$

Anticipating our results, let us call a task (S, \mathcal{P}) *feasible* if S and \mathcal{P} satisfy (8).

Let us take a moment to analyse condition (8). Note first that, like the given separations (A, B) and (C, D) , the new separation (E, F) will again distinguish P from P' : by assumption we have $(E, F) \in P$, and by (2) we have $(F, E) \leq (B \cap D, A \cup C) \leq (B, A) \in P'$, so $(F, E) \in P'$ by the fact that P' orients all of S consistently.

Now the idea behind (8) is that in our search for N we may find ourselves facing a choice between two crossing separations $(A, B), (C, D) \in S$ that both distinguish two profiles $P, P' \in \mathcal{P}$, and wonder which of these we should pick for N . (Clearly we cannot take both.) If (8) holds, we have the option to choose neither and pick (E, F) instead: it will do the job of distinguishing P from P' , and since it is nested with both (A, B) and (C, D) , putting it in N entails no prejudice to any future inclusion of either (A, B) or (C, D) in N .

Separations in S that do not distinguish any profiles in \mathcal{P} are not really needed for N , and so we may delete them.¹¹ So let us call a separation \mathcal{P} -

¹¹But do not have to: the freedom to discard or keep such separations will be our source of diversity for the tree-decompositions sought – which, as pointed out earlier, we may wish to endow with other desired properties than the minimum requirement of distinguishing \mathcal{P} .

relevant if it distinguishes some pair of profiles in \mathcal{P} , denote by R the set of all \mathcal{P} -relevant separations in S , and call (R, \mathcal{P}) the *reduction* of (S, \mathcal{P}) . If $(S, \mathcal{P}) = (R, \mathcal{P})$, we call this task *reduced*. Since all the separations (A, B) , (C, D) , (E, F) in (8) are \mathcal{P} -relevant, R inherits (8) from S (and vice versa):

$$(R, \mathcal{P}) \text{ is feasible if and only if } (S, \mathcal{P}) \text{ is feasible.} \quad (9)$$

Consider a task (S, \mathcal{P}) , fixed until Example 3 near the end of this section. Our aim is to construct N inductively, adding a few separations at each step. A potential danger when choosing a new separation to add to N is to pick one that crosses another separation that we might wish to include later. This can be avoided if we only ever add separations that are nested with all other separations in S that we might still want to include in N . So this will be our aim.

At first glance, this strategy might seem both wasteful and unrealistic: why should there even be a separation in S that we can choose at the start, one that is nested with all others? However, we cannot easily be more specific: since we want our nested subsystem N to be canonical, we are not allowed to break ties between crossing separations without appealing to an invariant of G as a criterion, and it would be hard to find such a criterion that applies to a large class of graphs without specifying this class in advance. But the strategy is also more realistic than it might seem. This is because the set of pairs of profiles we need to distinguish by separations still to be picked decreases as N grows. As a consequence, we shall need fewer separations in S to distinguish them. We may therefore be able to delete from S some separations that initially prevented the choice of a desired separation (A, B) for N by crossing it, because they are no longer needed to distinguish profiles in what remains of \mathcal{P} , thus freeing (A, B) for inclusion in N .

To get started, we thus have to look for separations (A, B) in S that are nested with all other separations in S . This will certainly be the case for (A, B) if, for every $(C, D) \in S$, we have either $(C, D) \leq (A, B)$ or $(D, C) \leq (A, B)$;¹² let us call such separations (A, B) *extremal* in S .

Consider distinct extremal separations (A, B) and (C, D) in S . Since \leq is a partial ordering, they cannot satisfy $(C, D) \leq (A, B) \leq (C, D)$. Hence, up to renaming, $(D, C) \leq (A, B)$, which shows that $\{(A, B), (C, D)\}$ is inconsistent:

$$\text{Distinct extremal separations are inconsistent. In particular, they} \quad (10) \\ \text{cannot lie in the same profile.}$$

If $(D, C) \leq (A, B)$, as above, then $(A, B) \not\leq (C, D)$, because $(D, C) \leq (C, D)$ would imply that $D \subseteq C$ and hence that (C, D) is improper (contrary to our assumption about S). But $(D, C) \leq (A, B)$ also implies that $(C, D) \not\leq (A, B)$, as otherwise $V = D \cup C \subseteq A$, and (A, B) would be improper. Hence,

$$\text{Distinct extremal separations are } \leq\text{-incomparable.} \quad (11)$$

Extremal separations always exist in a feasible task (S, \mathcal{P}) , as long as S contains no superfluous separations (which might cross useful ones):

¹²This implies that (A, B) is maximal in S , but only because we are assuming that all separations in S are proper: improper separations (C, D) can satisfy $(D, C) < (A, B) < (C, D)$.

Lemma 3.1. *If (S, \mathcal{P}) is feasible and reduced, then every \leq -maximal element of S is extremal in S .*

Proof. Let (A, B) be a maximal separation in S , and let $(C, D) \in S$ be any other separation. If (A, B) is nested with (C, D) it is comparable with (C, D) or (D, C) . Hence either $(C, D) \leq (A, B)$ or $(D, C) \leq (A, B)$ by the maximality of (A, B) , as desired. We may thus assume that (A, B) and (C, D) cross.

Since (S, \mathcal{P}) is reduced, (A, B) and (C, D) each distinguish two profiles from \mathcal{P} . Pick $P \in \mathcal{P}$ containing (A, B) . Since P orients all of S , it also contains (C, D) or (D, C) ; we assume it contains (C, D) . Now pick $P' \in \mathcal{P}$ containing (D, C) . If also $(B, A) \in P'$, then by (8) there exists an $(E, F) \in S \cap P$ such that $(A, B) \leq (A \cup C, B \cap D) \leq (E, F)$. Since (A, B) and (C, D) cross, the first of these inequalities is strict, which contradicts the maximality of (A, B) . Hence $(A, B) \in P' \cap P$. Since \mathcal{P} is reduced, there exists $P'' \in \mathcal{P}$ containing (B, A) . Applying (8) to P'' and either P or P' , we again find an $(E, F) > (A, B)$ that contradicts the maximality of (A, B) . \square

Note that the proof of Lemma 3.1 uses crucially that (S, \mathcal{P}) is feasible.

Lemma 3.2. *If (S, \mathcal{P}) is reduced, then for every extremal separation (A, B) in S there is a unique profile $P_{(A,B)} \in \mathcal{P}$ such that $(A, B) \in P_{(A,B)}$.*

Proof. As (S, \mathcal{P}) is reduced, there is a profile $P \in \mathcal{P}$ containing (A, B) . Suppose there is another such profile $P' \in \mathcal{P}$. Then P and P' are distinguished by some $(C, D) \in S$, because (S, \mathcal{P}) is a task. Since (A, B) is extremal, we may assume that $(C, D) \leq (A, B)$. Then the fact that (D, C) lies in one of P, P' contradicts the consistency of this profile. \square

By Lemma 3.2, (A, B) distinguishes $P_{(A,B)}$ from every other profile in \mathcal{P} .

Let us call a profile P orienting S *extremal* with respect to S if it contains an extremal separation of S . This will be the greatest, and hence the only maximal, separation in $P \cap S$.

As we have seen, an extremal profile is distinguished from every other profile in \mathcal{P} by some separation (A, B) that is nested with all the other separations in S ; this makes (A, B) a good choice for N . The fact that made (A, B) extremal, and hence nested with all other separations in S , was its maximality in S (Lemma 3.1). In the same way we may ask whether, given any profile $P \in \mathcal{P}$ (not necessarily extremal), the separations that are \leq -maximal in $P \cap S$ will be nested with every other separation in S : these are the separations ‘closest to P ’, much as (A, B) was closest to $P_{(A,B)}$ (although there can now be many such separations).

Let us prove that the following profiles have this property:

Definition 3.3. Call a profile P orienting S *well separated* in S if the set of \leq -maximal separations in $P \cap S$ is nested.

Note that extremal profiles are well separated.

Lemma 3.4. *Given a profile P orienting S , the following are equivalent:*

- (i) P is well separated in S .
- (ii) Every maximal separation in $P \cap S$ is nested with all of S .
- (iii) For every two crossing separations $(A, B), (C, D) \in P \cap S$ there exists a separation $(E, F) \in P \cap S$ such that $(A \cup C, B \cap D) \leq (E, F)$.

Proof. The implication (ii) \rightarrow (i) is trivial; we show (i) \rightarrow (iii) \rightarrow (ii).

(i) \rightarrow (iii): Suppose that P is well separated, and consider two crossing separations $(A, B), (C, D) \in P \cap S$. Let $(A', B') \geq (A, B)$ be maximal in $P \cap S$. Suppose first that $(A', B') \parallel (C, D)$. This means that (A', B') is \leq -comparable with either (C, D) or (D, C) . Since (A, B) is not nested with (C, D) we have $(A', B') \not\leq (C, D)$ and $(A', B') \not\leq (D, C)$, and since both (C, D) and (A', B') are in P , its consistency yields $(D, C) \not\leq (A', B')$. Hence $(C, D) \leq (A', B')$, and thus $(A \cup C, B \cap D) \leq (A', B')$. This proves (iii) with $(E, F) := (A', B')$.

Suppose now that (A', B') crosses (C, D) . Let $(C', D') \geq (C, D)$ be maximal in $P \cap S$. Since (A', B') and (C', D') are both maximal in $P \cap S$ they are nested, by assumption in (i). As in the last paragraph, now with (C', D') taking the role of (A', B') , and (A', B') taking the role of (C, D) , we can show that $(A, B) \leq (A', B') \leq (C', D')$ and hence $(A \cup C, B \cap D) \leq (C', D')$.¹³ This proves (iii) with $(E, F) := (C', D')$.

(iii) \rightarrow (ii): Suppose some maximal (A, B) in $P \cap S$ crosses some $(C, D) \in S$. As P orients S , and by symmetry of nestedness, we may assume that $(C, D) \in P$. By (iii), there is an $(E, F) \in P \cap S$ such that $(A \cup C, B \cap D) \leq (E, F)$, so $(A, B) \leq (E, F)$ as well as $(C, D) \leq (E, F)$. But then $(E, F) = (A, B)$ by the maximality of (A, B) , and hence $(A, B) \parallel (C, D)$, contradicting the choice of (A, B) and (C, D) . \square

Let us call a separation (A, B) *locally maximal* in (S, \mathcal{P}) if there exists a well-separated profile $P \in \mathcal{P}$ such that (A, B) is \leq -maximal in $P \cap S$. Lemma 3.4 shows that these separations are a good choice for inclusion in N :

Corollary 3.5. *Locally maximal separations in (S, \mathcal{P}) are nested with all of S .*

We have seen three ways of starting the construction of our desired nested subsystem $N \subseteq S$ by choosing for N some separations from S that are nested with all other separations in S : we may choose either

- the set $\text{ext}(S, \mathcal{P})$ of extremal separations in S and their inverses; or
- the set $\text{loc}(S, \mathcal{P})$ of all locally maximal separations in (S, \mathcal{P}) and their inverses; or
- the set $\text{all}(S, \mathcal{P})$ of all separations in S that are nested with every separation in S (which is a symmetric set).

¹³In the argument we need that (C, D) and (A', B') cross. This is why we first treated the case that they don't (but in that case we used that (A, B) and (C, D) cross).

Clearly,

$$\text{ext}(S, \mathcal{P}) \subseteq \text{all}(S, \mathcal{P}) \supseteq \text{loc}(S, \mathcal{P}) \quad (12)$$

in general, and

$$\emptyset \neq \text{ext}(S, \mathcal{P}) \subseteq \text{loc}(S, \mathcal{P}) \subseteq \text{all}(S, \mathcal{P}) \quad (13)$$

if $S \neq \emptyset$ and (S, \mathcal{P}) is feasible and reduced,¹⁴ since in that case every maximal separation in S is extremal (Lemma 3.1) and every extremal separation (A, B) is locally maximal for $P_{(A,B)} \in \mathcal{P}$.

Example 1. Let G consist of three large complete graphs X_1, X_2, X_3 threaded on a long path, as shown in Figure 1. Let S be the set of all proper 1-separations. Let $\mathcal{P} = \{P_1, P_2, P_3\}$, where P_i is the 2-profile induced by X_i . Then $\text{all}(S, \mathcal{P}) = S$, while $\text{loc}(S, \mathcal{P})$ contains only the separations in S with separators x_1, x_2, y_2 and x_3 , and $\text{ext}(S, \mathcal{P})$ only those with separator x_1 or x_3 .




Figure 1: Different results for $\text{ext}(S, \mathcal{P})$, $\text{loc}(S, \mathcal{P})$ and $\text{all}(S, \mathcal{P})$

How shall we proceed now, having completed the first step of our algorithm by specifying some nested subsystem $N \in \{\text{ext}(S, \mathcal{P}), \text{loc}(S, \mathcal{P}), \text{all}(S, \mathcal{P})\}$ of S ? The idea is that N divides G into chunks, which we now want to cut up further by adding more separations of S to N . While it is tempting to think of those ‘chunks’ as the N -blocks of G , it turned out that this fails to capture some of the more subtle scenarios. Here is an example:




Figure 2: Two S -distinguishable profiles living in an S -inseparable N -block

Example 2. Let G be the graph of Figure 2. Let N consist of the separations $(X_1, Y_1), \dots, (X_4, Y_4)$ and their inverses (Y_i, X_i) , where $Y_i := (A \cap B) \cup \bigcup_{j \neq i} X_j$, and let $S := N \cup \{(A, B), (B, A)\}$. Let \mathcal{P} consist of the following six profiles: the orientations of S towards X_1, \dots, X_4 , respectively, and two further profiles

¹⁴In fact, all we need for an extremal separation (A, B) to be locally maximal in a feasible (S, \mathcal{P}) is that it lies in some $P \in \mathcal{P}$. But this need not be the case if (S, \mathcal{P}) is not reduced: although one of (A, B) and (B, A) must lie in every $P \in \mathcal{P}$ (because P orients S), it might happen that this is always (B, A) .

P and P' which both orient N towards $A \cap B$ but of which P contains (A, B) while P' contains (B, A) . Then N distinguishes all these profiles except P and P' . But these are distinguished by (A, B) and (B, A) , so we wish to add these separations to N .

The profiles P and P' live in the same N -block of G , the set $A \cap B$. But although S distinguishes P from P' , it does not separate this N -block. We therefore cannot extend N to a separation system distinguishing \mathcal{P} by adding only separations from S that separate an N -block of G .

The lesson to be learnt from Example 2 is that the ‘chunks’ into which N divides our graph G should not be thought of as the N -blocks of G . An alternative that the example suggests would be to think of them as the N -blocks of \mathcal{P} : the equivalence classes of \mathcal{P} defined by how its profiles orient N . In the example, \mathcal{P} has five N -blocks: the four singleton N -blocks consisting just of the profile P_i that orients N towards X_i , and another N -block $\mathcal{Q} = \{P, P'\}$. So the algorithm could now focus on the subtask $(R_{\mathcal{Q}}, \mathcal{Q})$ with $R_{\mathcal{Q}} = \{(A, B), (B, A)\}$ consisting of those separations from S that distinguish profiles in \mathcal{Q} .

More generally, we could continue our algorithm after finding N by iterating it with the subtasks $(R_{\mathcal{Q}}, \mathcal{Q})$ of (S, \mathcal{P}) , where \mathcal{Q} runs over the non-trivial N -blocks of \mathcal{P} and $R_{\mathcal{Q}}$ is the set of \mathcal{Q} -relevant separations in S . This would indeed result in an overall algorithm that eventually produces a nested subsystem of S that distinguishes \mathcal{P} , solving our task (S, \mathcal{P}) .

However, when we considered our three alternative ways of obtaining N , we also had a secondary aim in mind: rather than working with the reduction (R, \mathcal{P}) of (S, \mathcal{P}) straight away, we kept our options open to include more separations in N than distinguishing \mathcal{P} requires, in order perhaps to produce a tree-decomposition into smaller parts.¹⁵ In the same spirit, our secondary aim now as we look for ways to continue our algorithm from N is not to exclude any separation of $S \setminus N$ from possible inclusion into N without need, i.e., to subdivide (S, \mathcal{P}) into subtasks (S_i, \mathcal{P}_i) if possible with $\bigcup_i S_i = S$.

In view of these two aims, the best way to think of the chunks left by N turned out to be neither as the (large) N -blocks of G , nor as the N -blocks of \mathcal{P} , but as something between the two: as the set \mathcal{O}_N of all consistent orientations of N . Let us look at these in more detail.

Recall that since every $P \in \mathcal{P}$ orients N , it defines an N -profile $P \cap N$. Equivalent P, P' define the same N -profile $P \cap N = P' \cap N$, the N -profile of the N -block \mathcal{Q} containing them. This is a consistent orientation of N . Conversely, given $O \in \mathcal{O}_N$, let us write \mathcal{P}_O for the set of profiles $P \in \mathcal{P}$ with $P \cap N = O$. Note that \mathcal{O}_N may also contain consistent orientations O of N , including N -profiles, that are not induced by any $P \in \mathcal{P}$, i.e., for which $\mathcal{P}_O = \emptyset$ (Fig. 1).

Similarly, every large N -block X of G defines an N -profile, the N -profile

¹⁵In Example 1 with $\text{ext}(S, \mathcal{P})$, where N consists of the proper 1-separations with separator x_1 or x_3 , every N -block of \mathcal{P} is trivial. But the middle N -block of G consists of X_2 and the entire path P , so we might cut it up further using the remaining 1-separations in S . If \mathcal{P} consisted only of P_1 and P_3 , then $\text{ext}(S, \mathcal{P})$ would have produced the same N , and the middle N -block would not even have a profile from \mathcal{P} living in it. But still, we might want to cut it up further.

$P_N(X)$ of X . This is a consistent orientation of N . Again, \mathcal{O}_N may also contain consistent orientations that are not of this form.¹⁶

Recall that a separation (A, B) splits $O \in \mathcal{O}_N$ if both $O \cup \{(A, B)\}$ and $O \cup \{(B, A)\}$ are again consistent.¹⁷ Let us write S_O for the set of separations in S that split O . These sets S_O extend our earlier sets R_Q in a way that encompasses all of $S \setminus N$, as intended:

Lemma 3.6. *Let N be a nested separation system that is oriented by every profile in \mathcal{P} and is nested with S .¹⁸*

- (i) $(S_O \mid O \in \mathcal{O}_N)$ is a partition of $S \setminus N$ (with $S_O = \emptyset$ allowed).
- (ii) $(\mathcal{P}_O \mid O \in \mathcal{O}_N)$ is a partition of \mathcal{P} (with $\mathcal{P}_O = \emptyset$ allowed).
- (iii) The N -profile P of any N -block Q of \mathcal{P} satisfies $\mathcal{P}_P = Q$ and $S_P \supseteq R_Q$.
- (iv) If (S, \mathcal{P}) is feasible, then the (S_O, \mathcal{P}_O) are feasible tasks.

Proof. (i) By Lemma 2.3, every separation $(A, B) \in S \setminus N$ splits a unique consistent orientation of N . Note that (A, B) is proper because (S, \mathcal{P}) is a task.

(ii) follows from the fact that every profile in \mathcal{P} orients N consistently.

(iii) The first assertion is immediate from the definition of an N -block of \mathcal{P} . For the second assertion let $(A, B) \in R_Q$ be given, distinguishing $Q, Q' \in \mathcal{Q}$ say. By (i), we have $(A, B) \in S_O$ for some $O \in \mathcal{O}_N$. Since Q and Q' are consistent, agree with P on N , and orient $\{(A, B), (B, A)\}$ differently, (A, B) splits the consistent orientation P of N . By the uniqueness of O this implies $P = O$. Hence, $(A, B) \in S_O = S_P$ as desired.

(iv) As S_O distinguishes \mathcal{P}_O , by (iii), we only have to show that (S_O, \mathcal{P}_O) is feasible. As (S, \mathcal{P}) is feasible, there is a separation (E, F) in S for any two crossing separations $(A, B), (C, D) \in S_O$ distinguishing profiles $P, P' \in \mathcal{P}_O$ as in (8). Since (E, F) also distinguishes P from P' , we have $(E, F) \in S_O$ by (iii). \square

We remark that the inclusion in Lemma 3.6 (iii) can be strict, since S_O may contain separations that do not distinguish any profiles in \mathcal{P} . Similarly, we can have $S_O \neq \emptyset$ for $O \in \mathcal{O}_N$ with $\mathcal{P}_O = \emptyset$.

The subtasks (S_O, \mathcal{P}_O) will be ‘easier’ than the original task (S, \mathcal{P}) , because we can reduce them further:

¹⁶In Example 2, the set $A \cap B$ is a small S -block of G for the nested separation system S . The profiles P, P' are two consistent orientations of S orienting it towards $A \cap B$, but not towards any large S -block.

¹⁷In Example 2, the N -profile of $X = A \cap B$ could be split into the consistent sets P and P' by adding the separations (A, B) and (B, A) , although the large N -block X could not be separated by any separation in S . Thus, splitting the N -profile of a large N -block is more subtle than separating the N -block itself.

Although all the consistent sets of separations considered in this example are in fact profiles, our aim to retain all the separations from $S \setminus N$ at this state requires that we do not restrict \mathcal{O}_N to profiles: there may be separations in S (which we want to keep) that only split a consistent orientation of N that is not a profile, or separations that split an N -profile into two consistent separations that are not profiles.

¹⁸For better applicability of the lemma later, we do not require that $N \subseteq S$.

Example 3. The separations (X, Y) and (X', Y') in Figure 3 are \mathcal{P} -relevant (because they separate the profiles $P, P' \in \mathcal{P}$, say), so they will not be deleted when we reduce S (which is, in fact reduced already). They both belong to S_O for the middle consistent orientation O of N , but are no longer \mathcal{P}_O -relevant, where $\mathcal{P}_O = \{P_1, P_2, P_3\}$ as shown. We can therefore discard them when we reduce the subtask (S_O, \mathcal{P}_O) before reapplying the algorithm to it, freeing (A, B) and (C, D) for adoption into N in the second step.




Figure 3: (X, Y) and (X', Y') are \mathcal{P} -relevant but no longer \mathcal{P}_O -relevant

More generally, reducing a subtask (S', \mathcal{P}') will be the crucial step in getting our algorithm back afloat if it finds no separation in S' that is nested with all the others. Example 3 shows that this can indeed happen.¹⁹ But after reducing (S', \mathcal{P}') to (R', \mathcal{P}') , say, we know from (13) that each of *ext*, *loc*, *all* will find a separation in R' that is nested with all the others.

As notation for the double step of first reducing a task (S, \mathcal{P}) , to (R, \mathcal{P}) say, and then applying *ext*, *loc* or *all*, let us define²⁰

$$\text{ext}_r(S, \mathcal{P}) := \text{ext}(R, \mathcal{P}); \quad \text{loc}_r(S, \mathcal{P}) := \text{loc}(R, \mathcal{P}); \quad \text{all}_r(S, \mathcal{P}) := \text{all}(R, \mathcal{P}).$$

We shall view each of *ext*, *loc*, *all*, *ext_r*, *loc_r*, *all_r* as a function that maps a given graph G and a feasible task (S, \mathcal{P}) in G to a nested subsystem N' of S' .

A *strategy* is a map $\sigma: \mathbb{N} \rightarrow \{\text{ext}, \text{loc}, \text{all}, \text{ext}_r, \text{loc}_r, \text{all}_r\}$ such that $\sigma(i) \in \{\text{ext}_r, \text{loc}_r, \text{all}_r\}$ for infinitely many i . The idea is that, starting from some feasible task (S, \mathcal{P}) , we apply $\sigma(i)$ at the i th step of the algorithm to the subtasks produced by the previous step, adding more and more separations to N . The requirement that for infinitely many i we have to reduce the subtasks first ensures that we cannot get stuck before N distinguishes all of \mathcal{P} .

Formally, we define a map $(\sigma, G, (S, \mathcal{P})) \mapsto N_\sigma(S, \mathcal{P})$ by which every strategy σ *determines* for every feasible task (S, \mathcal{P}) in a graph G some set $N_\sigma(S, \mathcal{P})$. We define this map recursively, as follows. Define σ^+ by setting $\sigma^+(i) := \sigma(i+1)$

¹⁹More generally, if we apply *all* (S, \mathcal{P}) in the first step to obtain N , say, then every subtask (S_O, \mathcal{P}_O) with $O \in N_O$ will have this property: if a separation $(A, B) \in S_O$ was nested with all of S_O it would in fact be nested with all of S (and have been included in N), by Lemma 2.3 (ii).

²⁰For the remainder of this section, G and (S, \mathcal{P}) will no longer be fixed.

for all $i \in \mathbb{N}$. Note that if σ is a strategy then so is σ^+ . Let $s := |S|$, and let r_σ be the least integer r such that $\sigma(r) \in \{\text{ext}_r, \text{loc}_r, \text{all}_r\}$. Our recursion is on s , and for fixed s on r_σ , for all G .

If $s = 0$, we let $N_\sigma(S, \mathcal{P}) := S = \emptyset$. Suppose now that $s \geq 1$; thus, $S \neq \emptyset$. Let $N := \sigma(0)(S, \mathcal{P})$. By Lemma 3.6 (iv), the subtasks (S_O, \mathcal{P}_O) with $O \in \mathcal{O}_N$ are again feasible.

Assume first that $r_\sigma = 0$, i.e. that $\sigma(0) \in \{\text{ext}_r, \text{loc}_r, \text{all}_r\}$, and let (R, \mathcal{P}) be the reduction of (S, \mathcal{P}) . If $R \subsetneq S$ we let $N_\sigma(S, \mathcal{P}) := N_\sigma(R, \mathcal{P})$, which is already defined. If $R = S$ then $N \neq \emptyset$ by (13), and $|S_O| \leq |S \setminus N| < s$ for every $O \in \mathcal{O}_N$. Thus, $N_{\sigma^+}(S_O, \mathcal{P}_O)$ is already defined.

Assume now that $r_\sigma > 0$, i.e. that $\sigma(0) \in \{\text{ext}, \text{loc}, \text{all}\}$. Then $r_{\sigma^+} < r_\sigma$, so again $N_{\sigma^+}(S_O, \mathcal{P}_O)$ is already defined. In either case we let

$$N_\sigma(S, \mathcal{P}) := N \cup \bigcup_{O \in \mathcal{O}_N} N_{\sigma^+}(S_O, \mathcal{P}_O). \quad (14)$$

Theorem 3.7. *Every strategy σ determines for every feasible task (S, \mathcal{P}) in a graph G a nested subsystem N_σ of S that distinguishes all the profiles in \mathcal{P} .*

These sets N_σ are canonical in that, for each σ , the map $(G, S, \mathcal{P}) \mapsto N_\sigma$ commutes with all isomorphisms $G \mapsto G'$. In particular, if S and \mathcal{P} are invariant under the automorphisms of G , then so is N_σ .

Proof. We apply induction along the recursion in the definition of $N_\sigma = N_\sigma(S, \mathcal{P})$. If $s = 0$, then $N_\sigma = S$ distinguishes all the profiles in \mathcal{P} , because (S, \mathcal{P}) is a task.

Suppose now that $s \geq 1$. Then N_σ is defined by (14). Both N and the sets $N_{\sigma^+}(S_O, \mathcal{P}_O)$ are subsets of S , hence so is N_σ . By definition, N is nested with all of S , in particular, with itself and the sets $N_{\sigma^+}(S_O, \mathcal{P}_O)$. These sets are themselves nested by induction, and nested with each other by Lemma 2.3 (ii). Thus, N_σ is a nested subset of S .

Any two profiles in the same N -block of \mathcal{P} are, by induction, distinguished by $N_{\sigma^+}(S_O, \mathcal{P}_O)$ for their common consistent orientation O (cf. Lemma 3.6 (iii)). Profiles from different N -blocks of \mathcal{P} are distinguished by N . Hence N_σ distinguishes \mathcal{P} .

Finally, the maps $(S, \mathcal{P}) \mapsto N_\sigma$ commute with all isomorphisms $G \mapsto G'$. Indeed, the maps $(S, \mathcal{P}) \mapsto N$ and hence $(S, \mathcal{P}) \mapsto \{(S_O, \mathcal{P}_O) \mid O \in \mathcal{O}_N\}$ do by definition of $\text{ext}, \text{loc}, \text{all}, \text{ext}_r, \text{loc}_r, \text{all}_r$, and the maps $(S_O, \mathcal{P}_O) \mapsto N_{\sigma^+}(S_O, \mathcal{P}_O)$ do by induction. \square

Let us complete this section with an example of how the use of different strategies can yield different nested separation systems. Unlike in the simpler Example 1, these will not extend each other, but will be incomparable under set inclusion. Let Ext , Loc and All denote the strategies given by setting $\text{Ext}(i) = \text{ext}_r$ and $\text{Loc}(i) = \text{loc}_r$ and $\text{All}(i) = \text{all}_r$, respectively, for all $i \in \mathbb{N}$.

Example 4. Let G be the 3-connected graph obtained from a (3×17) -grid by attaching two K^4 s at its short ends, and some further edges as in Figure 4. Let S be the set of all its proper 3-separations, and \mathcal{P} the set of all its 4-block




Figure 4: Three different nested separation systems distinguishing the 4-blocks profiles. It is not hard to show, and will follow from Lemma 4.1, that (S, \mathcal{P}) is a feasible task.

The grey bars in each of the three copies of the graph highlight the separators of the separations in $N_{\text{Ext}}(S, \mathcal{P})$, in $N_{\text{Loc}}(S, \mathcal{P})$, and in $N_{\text{All}}(S, \mathcal{P})$. The step at which a separator was added is indicated by a number.

Note that the three nested separation systems obtained are not only \subseteq -incomparable. They are not even nested with each other: for every pair of N_{Ext} , N_{Loc} and N_{All} we can find a pair of crossing separations, one from either system.

4 Iterated strategies and tree-decompositions

Let us apply the results of Section 3 to our original problem of finding, for any set \mathcal{P} of k -profiles in G , within the set S of all proper ($< k$)-separations of G a nested subset that distinguishes \mathcal{P} (and hence gives rise to a tree-decomposition of adhesion $< k$ that does the same). This is easy now if G is $(k-1)$ -connected:

Lemma 4.1. *If G is $(k-1)$ -connected ($k \geq 1$), then (S, \mathcal{P}) is a feasible task.*

Proof. By (4), the pair (S, \mathcal{P}) is a task: since all k -profiles contain the same improper separations, the separations (A, V) with $|A| < k$, distinct k -profiles must differ on S and are thus distinguished by S . So let us show that this

task is feasible (8). Let (A, B) and (C, D) be crossing separations in S , and let $P, P' \in \mathcal{P}$ be such that $(A, B), (C, D) \in P$ and $(B, A), (D, C) \in P'$. We first prove that $(E, F) := (A \cup C, B \cap D)$ has order at most $k - 1$.

Suppose (E, F) has order greater than $k - 1$. By Lemma 2.1 this implies that the separation $(X, Y) := (B \cup D, A \cap C)$ has order less than $k - 1$, and hence is improper since G is $(k - 1)$ -connected. As both (A, B) and (C, D) are proper separations and hence $X \setminus Y = (B \cup D) \setminus (A \cap C) \neq \emptyset$, we then have $Y \subseteq X$. Then $(X, Y) \notin P'$, by (4). But by definition of (X, Y) and (P) for P' we also have $(Y, X) \notin P'$. This contradicts (5).

We have shown that (E, F) has order at most $k - 1$. By (5) and (P) this implies that $(E, F) \in P$. To complete our proof of (8), it remains to show that $(E, F) \in S$, i.e. that (E, F) is proper. If it is improper then $F \subseteq E = V$, since $E \setminus F \supseteq A \setminus B \neq \emptyset$ and therefore $F \neq V$. By (4), this contradicts the fact that $(E, F) \in P$. \square

Coupled with Lemma 4.1, we can apply Theorem 3.7 as follows:

Corollary 4.2. *Every strategy σ determines for every $(k - 1)$ -connected graph G a canonical nested system of separations of order $k - 1$ which distinguishes all the k -profiles in G .* \square

If G is not $(k - 1)$ -connected, the task (S, \mathcal{P}) consisting of the set S of all proper $(< k)$ -separations of G and the given set \mathcal{P} of k -profiles need not be feasible. Indeed, the separation $(A \cup C, B \cap D)$ in (8) might have order $\geq k$ even if both (A, B) and (C, D) have order $< k$. Then if $B \cap D$ induces a big complete graph, for example, there will be no (E, F) as required in (8).

However, if $|A \cap B| = |C \cap D| = k - 1$ in this example, the separation $(B \cup D, A \cap C)$ will have some order $\ell < k - 1$. This separation, too, distinguishes the profiles P, P' given in (8), and thus by (7) lies in S . Hence our dilemma of having to choose between (A, B) and (C, D) for inclusion in our nested subset N of S (which gave rise to (8) and the notion of feasibility) would not occur if we considered lower-order separations first: we would then have included $(B \cup D, A \cap C)$ in N , and would need neither (A, B) nor (C, D) to distinguish P from P' .

It turns out that this approach does indeed work in general. Given our set \mathcal{P} of k -profiles, let us define for any $1 \leq \ell \leq k$ and $P \in \mathcal{P}$ the induced ℓ -profile

$$P_\ell := \{(A, B) \in P : |A \cap B| < \ell\}, \quad \text{and set } \mathcal{P}_\ell := \{P_\ell \mid P \in \mathcal{P}\}.$$

Note that distinct k -profiles P may induce the same ℓ -profile P_ℓ . Let $\kappa(P, P')$ denote the least order of any separation in G that distinguishes two profiles P, P' .

The idea now is to start with a nested set $N_1 \subseteq S$ of (< 1) -separations that distinguishes \mathcal{P}_1 , then to extend N_1 to a set $N_2 \subseteq S$ of (< 2) -separations that distinguishes \mathcal{P}_2 , and so on. The tasks (S_O, \mathcal{P}_O) to be solved at step k , those left by the consistent orientations O of N_{k-1} , will then be feasible: since N_{k-1} already distinguishes \mathcal{P}_{k-1} , and hence distinguishes any $P, P' \in \mathcal{P}$ with $\kappa(P, P') < k - 1$, any P, P' in a common \mathcal{P}_O will satisfy $\kappa(P, P') = k - 1$, and (8) will follow from Lemma 2.1 as in the proof of Lemma 4.1.

What is harder to show is that those (S_O, \mathcal{P}_O) are indeed tasks: that S_O is rich enough to distinguish \mathcal{P}_O . This will be our next lemma. Let us say that a separation (A, B) of order ℓ that distinguishes two profiles P and P' does so *efficiently* if $\kappa(P, P') = \ell$. We say that (A, B) is \mathcal{P} -*essential* if it efficiently distinguishes some pair of profiles in \mathcal{P} . Note that for $\ell \leq m$ we have $(\mathcal{P}_m)_\ell = \mathcal{P}_\ell$, and if (A, B) is \mathcal{P}_ℓ -essential it is also \mathcal{P}_m -essential.

Lemma 4.3. *Let \mathcal{P} be a set of k -profiles in G , let N be a nested system of \mathcal{P}_{k-1} -essential separations of G that distinguishes all the profiles in \mathcal{P}_{k-1} efficiently, and let S be the set of all proper $(k-1)$ -separations of G that are nested with N . Then for every consistent orientation O of N the pair (S_O, \mathcal{P}_O) is a feasible task.*

Proof. As pointed out earlier, (S_O, \mathcal{P}_O) will clearly be feasible once we know it is a task. Since all profiles in \mathcal{P}_O are k -profiles and hence orient S_O , we only have to show that S_O distinguishes \mathcal{P}_O .

So consider distinct profiles $P_1, P_2 \in \mathcal{P}_O$. Being k -profiles, they are distinguished by a proper separation (A, B) of order at most $k-1$. Choose (A, B) nested with as many separations in N as possible; we shall prove that it is nested with all of N , giving $(A, B) \in S_O$ as desired. Note that $|A \cap B| = k-1$, since N does not distinguish $P_1, P_2 \in \mathcal{P}_O$. As (A, B) distinguishes P_1 from P_2 , we may assume $(B, A) \in P_1$ and $(A, B) \in P_2$.

Suppose (A, B) crosses a separation $(C, D) \in N$. Since every separation in N is \mathcal{P}_{k-1} -essential, by assumption, there are profiles $Q'_1, Q'_2 \in \mathcal{P}_{k-1}$ such that (C, D) distinguishes Q'_1 from Q'_2 efficiently. By definition of \mathcal{P}_{k-1} , this implies that there are distinct profiles $Q_1, Q_2 \in \mathcal{P}$ which (C, D) distinguishes efficiently. Then

$$m := |C \cap D| < k-1 = |A \cap B|. \quad (15)$$

Hence (C, D) does not distinguish P_1 from P_2 ; we assume that $(D, C) \in P_1 \cap P_2 \cap Q_1$ and $(C, D) \in Q_2$. Since Q_2 is a k -profile it contains precisely one of (A, B) and (B, A) , we assume $(A, B) \in Q_2$ (Figure 5).




Figure 5: The known positions of P_1, P_2, Q_1 and Q_2

Now if $(X, Y) := (A \cup C, B \cap D)$ has order $< m$, then $(X, Y) \in Q_2$ by (P), and $(Y, X) \in Q_1$ by consistency. Hence (X, Y) distinguishes Q_1 from Q_2 and has smaller order than (C, D) , contradicting the fact that (C, D) distinguishes Q_1 and Q_2 efficiently. Thus (X, Y) has order at least m .

Hence by (15) and Lemma 2.1, the order of $(E, F) := (B \cup D, A \cap C)$ is at most $k - 1$. Then $(E, F) \in P_1$ by (P), and $(F, E) \in P_2$ by consistency. Thus (E, F) distinguishes P_1 from P_2 . By [3, Lemma 2.2],²¹ (E, F) is nested with every separation in N that (A, B) is nested with, and in addition (E, F) is also nested with (C, D) . Hence, (E, F) is nested with more separations in N than (A, B) is, contradicting the choice of (A, B) . \square

When we apply Lemma 4.3 inductively, we have to make sure that every N_ℓ we construct consists only of \mathcal{P}_ℓ -essential separations. To ensure this, we have to reduce any task we tackle in the process of constructing N_ℓ . Given $k \geq 1$, a k -strategy is a k -tuple $(\sigma_1, \dots, \sigma_k)$ of strategies σ_i each with range $\{\text{ext}_r, \text{loc}_r, \text{all}_r\}$. The restriction in the range of k -strategies will reduce our freedom in shaping the decompositions, but Example 4 shows that considerable diversity remains.

Given G and \mathcal{P} , a k -strategy $\Sigma = (\sigma_1, \dots, \sigma_k)$ determines the set $N_\Sigma = N_\Sigma(G, \mathcal{P})$ defined recursively as follows. For $k = 1$, let $N_\Sigma := N_{\sigma_1}(S, \mathcal{P})$, where S is the set of proper (< 1)-separations of G . Then for $k \geq 2$ let

$$N_\Sigma := N \cup \bigcup_{O \in \mathcal{O}_N} N_{\sigma_k}(S_O, \mathcal{P}_O), \quad (16)$$

where $N = N_{\Sigma'}(G, \mathcal{P}_{k-1})$ for $\Sigma' = (\sigma_1, \dots, \sigma_{k-1})$, and S is the set of proper $(k-1)$ -separations of G that are nested with N . The pairs (S_O, \mathcal{P}_O) are defined with reference to N , S and $\mathcal{P} = \mathcal{P}_k$ as before Lemma 3.6.

As before, the sets N_Σ will be *canonical* in that, for each Σ , the map $(G, \mathcal{P}) \mapsto N_\Sigma$ commutes with all isomorphisms $G \mapsto G'$. In particular, if \mathcal{P} is invariant under the automorphisms of G , then so is N_Σ .

Theorem 4.4. *Every k -strategy Σ determines for every set \mathcal{P} of k -profiles in a graph G a canonical nested system $N_\Sigma(G, \mathcal{P})$ of \mathcal{P} -essential separations of order $< k$ that distinguishes all the profiles in \mathcal{P} efficiently.*

Proof. We show by induction on k that the recursive definition of $N_\Sigma(G, \mathcal{P})$ succeeds and that $N_\Sigma = N_\Sigma(G, \mathcal{P})$ has the desired properties. For $k = 1$ this follows from Corollary 4.2.

For $k \geq 2$ let N and S be defined as before the theorem. By the induction hypothesis, N is a nested system of \mathcal{P}_{k-1} -essential separations of G that distinguishes the profiles in \mathcal{P}_{k-1} efficiently. For every consistent orientation O of N the pair (S_O, \mathcal{P}_O) is a feasible task, by Lemma 4.3. By Theorem 3.7, then, σ_k determines a nested separation system $N_{\sigma_k}(S_O, \mathcal{P}_O) \subseteq S_O \subseteq S$ that distinguishes all the profiles in \mathcal{P}_O . By definition of S , all these $N_{\sigma_k}(S_O, \mathcal{P}_O)$ are nested with N , and they are nested with each other by Lemma 2.3. Hence N_Σ is well defined by (16) and forms a nested separation system.

²¹Swap the names of (C, D) and (E, F) in the statement of the lemma in [3].

Let us show that N_Σ has the desired properties. To show that N_Σ distinguishes the profiles in \mathcal{P} efficiently, consider distinct $P, Q \in \mathcal{P}$. If $\kappa(P, Q) < k-1$, then $P_{k-1} \neq Q_{k-1}$ are distinct profiles in \mathcal{P}_{k-1} . So by the induction hypothesis there is a separation in $N \subseteq N_\Sigma$ that distinguishes P from Q efficiently. If $\kappa(P, Q) = k-1$, we have $P_{k-1} = Q_{k-1}$. Then P and Q have the same N -profile O , and $P, Q \in \mathcal{P}_O$. Hence there is a separation in $N_{\sigma_k}(S_O, \mathcal{P}_O) \subseteq N_\Sigma$ that distinguishes P from Q ; as it has order $k-1$, it does so efficiently.

It remains to show that every separation $(A, B) \in N_\Sigma$ is \mathcal{P} -essential. If $(A, B) \in N$, this holds by the induction hypothesis and the definition of \mathcal{P}_{k-1} . So assume that $(A, B) \in N_\Sigma \setminus N$. Then there is a consistent orientation O of N such that $(A, B) \in N_{\sigma_k}(S_O, \mathcal{P}_O)$. Since $\sigma_k(i) \in \{\text{ext}_r, \text{loc}_r, \text{all}_r\}$ for all $i \in \mathbb{N}$, we know that (A, B) distinguishes some $P, Q \in \mathcal{P}_O$. Then $\kappa(P, Q) = k-1 = |A \cap B|$, as otherwise N would distinguish P from Q by the induction hypothesis. Hence, (A, B) distinguishes P from Q efficiently, as desired. \square

It remains to translate our results from separation systems to tree-decompositions. Recall from Theorem 2.2 that every nested proper separation system N of G is induced by some tree-decomposition (T, \mathcal{V}) : the separations of G that correspond to edges of T are precisely those in N . In [3] we showed that (T, \mathcal{V}) is uniquely determined by N .²² Hence if N is determined by some k -strategy, as in Theorem 4.4, we may say that this k -strategy *defines* (T, \mathcal{V}) on G .

If N comes from an application of Theorem 4.4, it will be proper (since it consists of essential separations (7)) and canonical. In particular, if the set \mathcal{P} of profiles considered is invariant under the automorphisms of G , then so is N , and hence so is T : the automorphisms of G will act on $V(T)$ as automorphisms of T . And many natural choices of \mathcal{P} are invariant under the automorphisms of G : the set of all k -profiles for given k , the set of all k -block profiles, or the set of all tangles of order k to name some examples. All these can thus be distinguished in a single tree-decomposition:

Theorem 4.5. *Given $k \in \mathbb{N}$ and a graph G , every k -strategy defines a canonical tree-decomposition of adhesion $< k$ of G that distinguishes all its k -blocks and tangles of order k . In particular, such a decomposition exists.* \square

Theorem 4.5 is not the end of this story, but rather a beginning. One can now build on the fact that these tree-decompositions are given constructively and study their details. For example, we may wish to find out more about the structure or size of their parts, or obtain bounds on the number of parts containing a k -block or accommodating a tangle of order k , compared with the total number of parts. The answers to such questions will depend on the k -strategy chosen. We shall pursue such questions in [1].

²²We assume here that parts corresponding to different nodes of T are distinct. It is always possible to artificially enlarge the tree without changing the set of separations by duplicating a part.

References

- [1] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs II. The parts. Preprint 2013.
- [2] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. k -Blocks: a connectivity invariant for graphs. Preprint 2013.
- [3] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. Connectivity and tree structure in finite graphs. *Combinatorica*, 34(1):1–35, 2014.
- [4] R. Diestel. *Graph Theory*. Springer, 4th edition, 2010.
- [5] M.J. Dunwoody and B. Krön. Vertex cuts. arXiv:0905.0064, 2009.
- [6] M. Grohe and D. Marx. Structure theorem and isomorphism test for graphs with excluded topological subgraphs (STOC'12). arXiv:1111.1109, 2011.
- [7] F. Hundertmark. Profiles. An algebraic approach to combinatorial connectivity. arXiv:1110.6207, 2011.
- [8] N. Robertson and P.D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)*, 52:153–190, 1991.
- [9] W. T. Tutte. *Graph Theory*. Addison-Wesley, 1984.