

Connectivity and tree structure in finite graphs

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

1 February, 2013

Abstract

Considering systems of separations in a graph that separate every pair of a given set of vertex sets that are themselves not separated by these separations, we determine conditions under which such a separation system contains a nested subsystem that still separates those sets and is invariant under the automorphisms of the graph.

As an application, we show that the k -blocks – the maximal vertex sets that cannot be separated by at most k vertices – of a graph G live in distinct parts of a suitable tree-decomposition of G of adhesion at most k , whose decomposition tree is invariant under the automorphisms of G . This extends recent work of Dunwoody and Krön and, like theirs, generalizes a similar theorem of Tutte for $k = 2$.

Under mild additional assumptions, which are necessary, our decompositions can be combined into one overall tree-decomposition that distinguishes, for all k simultaneously, all the k -blocks of a finite graph.

1 Introduction

Ever since graph connectivity began to be systematically studied, from about 1960 onwards, it has been an equally attractive and elusive quest to ‘decompose a k -connected graph into its $(k + 1)$ -connected components’. The idea was modelled on the well-known block-cutvertex tree, which for $k = 1$ displays the global structure of a connected graph ‘up to 2-connectedness’. For general k , the precise meaning of what those ‘ $(k + 1)$ -connected components’ should be varied, and came to be considered as part of the problem. But the aim was clear: it should allow for a decomposition of the graph into those ‘components’, so that their relative structure would display some more global structure of the graph.

While originally, perhaps, these ‘components’ were thought of as subgraphs, it soon became clear that, for larger k , they would have to be defined differently. For $k = 2$, Tutte [11] found a decomposition which, in modern terms,¹ would

¹Readers not acquainted with the terminology of graph minor theory can skip the details of this example without loss. The main point is that those ‘torsos’ are not subgraphs, but subgraphs plus some additional edges reflecting the additional connectivity that the rest of the graph provides for their vertices.

be described as a tree-decomposition of adhesion 2 whose torsos are either 3-connected or cycles.

For general k , Robertson and Seymour [10] re-interpreted those ‘ $(k + 1)$ -connected components’ in a radically new (but interesting) way as ‘tangles of order k ’. They showed, as a cornerstone of their theory on graph minors, that every finite graph admits a tree-decomposition that separates all its maximal tangles, regardless of their order, in that they inhabit different parts of the decomposition. Note that this solves the modified problem for all k simultaneously, a feature we shall achieve also for the original problem.

More recently still, Dunwoody and Krön [6], taking their lead directly from Tutte (and from Dunwoody’s earlier work on tree-structure induced by edge-cuts [5]), followed up Tutte’s observation that his result for $k = 2$ can alternatively be described as a tree-like decomposition of a graph G into cycles and vertex sets that are ‘2-inseparable’: such that no set of at most 2 vertices can separate any two vertices of that set in G . Note that such ‘ k -inseparable’ sets of vertices, which were first studied by Mader [8], differ markedly from k -connected subgraphs, in that their connectivity resides not on the set itself but in the ambient graph. For example, joining $r > k$ isolated vertices pairwise by $k + 1$ independent paths of length 2, all disjoint, makes this set into a ‘ k -block’, a maximal k -inseparable set of vertices. This then plays an important structural (hub-like) role for the connectivity of the graph, but it is still independent.

External connectivity of a set of vertices in the ambient graph had been considered before in the context of tree-decompositions and tangles [4, 9]. But it was Dunwoody and Krön who realized that k -inseparability can serve to extend Tutte’s result to $k > 2$: they showed that the k -blocks of a finite k -connected graph can, in principle, be separated canonically in a tree-like way [6]. We shall re-prove this in a simpler and stronger form, extend it to graphs of arbitrary connectivity, and cast the ‘tree-like way’ in the standard form of tree-decompositions. We show in particular that every finite graph has a canonical tree-decomposition of adhesion at most k such that distinct k -blocks are contained in different parts (Theorem 1); this appears to solve the original problem for fixed k in a strongest-possible way. For graphs whose k -blocks have size at least $3k/2$ for all k , a weak but necessary additional assumption, these decompositions can be combined into one unified tree-decomposition that distinguishes all the blocks of the graph, simultaneously for all k (Theorem 2).

Our paper is independent of the results stated in [6].² Our approach will be as follows. We first develop a more general theory of separation systems to deal with the following abstract problem. Let \mathcal{S} be a set of separations in a graph, and let \mathcal{I} be a collection of \mathcal{S} -inseparable sets of vertices, sets which, for every separation $(A, B) \in \mathcal{S}$, lie entirely in A or entirely in B . Under what condition does \mathcal{S} have a *nested* subsystem \mathcal{N} that still separates all the sets in \mathcal{I} ? In a further step we show how such nested separation systems \mathcal{N} can be captured by tree-decompositions.³

²The starting point for this paper was that, despite some effort, we were unable to verify some of the results claimed in [6].

³It is easy to see that tree-decompositions give rise to nested separation systems. The

The gain from having an abstract theory of how to extract nested subsystems from a given separation system is its flexibility. For example, we shall use it in [7] to prove that every finite graph has a canonical (in the sense above) tree-decomposition separating all its maximal tangles. This improves on the result of Robertson and Seymour [10] mentioned earlier, in that their decomposition is not canonical in our sense: it depends on an assumed vertex enumeration to break ties when choosing which of two crossing separations should be picked for the nested subsystem. The choices made by our decompositions will depend only on the structure of the graph. In particular, they will be invariant under its automorphisms, which thus act naturally also on the set of parts of the decomposition and on the associated decomposition tree.

To state our main results precisely, let us define their terms more formally. In addition to the terminology explained in [3] we say that a set X of vertices in a graph G is k -inseparable in G if $|X| > k$ and no set S of at most k vertices separates two vertices of $X \setminus S$ in G . A maximal k -inseparable set of vertices is a k -block,⁴ or simply a *block*. The smallest k for which a block is a k -block is the *rank* of that block; the largest such k is its *order*.

The intersections $V_t \cap V_{t'}$ of ‘adjacent’ parts in a tree-decomposition $(\mathcal{T}, \mathcal{V})$ of G (those for which tt' is an edge of \mathcal{T}) are the *adhesion sets* of $(\mathcal{T}, \mathcal{V})$; the maximum size of such a set is the *adhesion* of $(\mathcal{T}, \mathcal{V})$. A tree-decomposition of adhesion at most k *distinguishes* two k -blocks b_1, b_2 of G if they are contained in different parts, V_{t_1} and V_{t_2} say. It does so *efficiently* if the t_1 – t_2 path in the decomposition tree \mathcal{T} has an edge tt' whose adhesion set (which will separate b_1 from b_2 in G) has size $\kappa(b_1, b_2)$, the minimum size of a b_1 – b_2 separator in G . The tree-decomposition $(\mathcal{T}, \mathcal{V})$ is *Aut(G)-invariant* if the automorphisms of G act on the set of parts in a way that induces an action on the tree \mathcal{T} .

Theorem 1. *Given any integer $k \geq 0$, every finite graph G has an Aut(G)-invariant tree-decomposition of adhesion at most k that efficiently distinguishes all its k -blocks.*

Unlike in the original problem, the graph G in Theorem 1 is not required to be k -connected. This is a more substantial improvement than it might seem. It becomes possible only by an inductive approach which refines, for increasing $\ell = 0, 1, \dots$, each part of a given tree-decomposition of G of adhesion at most ℓ by a finer tree-decomposition of adhesion at most $\ell + 1$, until for $\ell = k$ the desired decomposition is achieved. The problem with this approach is that, in general, a graph G need not admit a unified tree-decomposition that distinguishes its ℓ -blocks for all $\ell \in \mathbb{N}$ simultaneously. Indeed, we shall see in Section 6 an example where G has two ℓ -blocks separated by a unique separation of order at most ℓ , as well as two $(\ell + 1)$ -blocks separated by a unique separation of order

converse is less clear.

⁴Belonging to the same k -block is not an equivalence relation on $V(G)$, but almost: distinct k -blocks can be separated by k or fewer vertices. A long cycle has exactly one k -block for $k \in \{0, 1\}$ and no k -block for $k \geq 2$. A large grid has a unique k -block for $k \in \{0, 1\}$, five 2-blocks (each of the corner vertices with its neighbours, plus the set of non-corner vertices), and one 3-block (the set of its inner vertices). It has no k -block for $k \geq 4$.

at most $\ell + 1$, but where these two separations ‘cross’: we cannot adopt both for the same tree-decomposition of G . The reason why this inductive approach nonetheless works for a proof of Theorem 1 is that we aim for slightly less there: at stage ℓ we only separate those ℓ -blocks of G that contain a k -block for the fixed k given in the theorem, not all the ℓ -blocks of G .

However, there is a slight strengthening of the notion of a block that does make it possible to construct an overall tree-decomposition separating all the blocks of a graph at once. We shall call such blocks *robust*. Their precise definition is technical and will be given later; it essentially describes the exact way in which the offending block of the above counterexample lies in the graph.⁵ In practice ‘most’ blocks of a graph will be robust, including all k -blocks that are complete or have size at least $3k/2$.

If all the blocks of a graph G are robust, how will they lie in the unified tree-decomposition of G that distinguishes them all? Some blocks (especially those of large order) will reside in a single part of this decomposition, while others (of smaller order) will inhabit a subtree consisting of several parts. Subtrees accommodating distinct k -blocks, however, will be disjoint. Hence for any fixed k we can contract them to single nodes, to reobtain the tree-decomposition from Theorem 1 in which the k -blocks (for this fixed k) inhabit distinct single parts. As k grows, we thus have a sequence $(\mathcal{T}_k, \mathcal{V}_k)_{k \in \mathbb{N}}$ of tree-decompositions, each refining the previous, that gives rise to our overall tree-decomposition in the last step of the sequence.

Formally, let us write $(\mathcal{T}_m, \mathcal{V}_m) \preceq (\mathcal{T}_n, \mathcal{V}_n)$ for tree-decompositions $(\mathcal{T}_m, \mathcal{V}_m)$ and $(\mathcal{T}_n, \mathcal{V}_n)$ if the decomposition tree \mathcal{T}_m of the first is a minor of the decomposition tree \mathcal{T}_n of the second, and a part $V_t \in \mathcal{V}_m$ of the first decomposition is the union of those parts $V_{t'}$ of the second whose nodes t' were contracted to the node t of \mathcal{T}_m .

Theorem 2. *For every finite graph G there is a sequence $(\mathcal{T}_k, \mathcal{V}_k)_{k \in \mathbb{N}}$ of tree-decompositions such that, for all k ,*

- (i) $(\mathcal{T}_k, \mathcal{V}_k)$ has adhesion at most k and distinguishes all robust k -blocks;
- (ii) $(\mathcal{T}_k, \mathcal{V}_k) \preceq (\mathcal{T}_{k+1}, \mathcal{V}_{k+1})$;
- (iii) $(\mathcal{T}_k, \mathcal{V}_k)$ is $\text{Aut}(G)$ -invariant.

The decomposition $(\mathcal{T}_k, \mathcal{V}_k)$ will in fact distinguish distinct robust k -blocks b_1, b_2 efficiently, by (i) for $k' = \kappa(b_1, b_2)$ and (ii). In Section 6 we shall prove Theorem 2 in a stronger form, which also describes how blocks of different rank or order are distinguished.

This paper is organized as follows. In Section 2 we collect together some properties of pairs of separations, either crossing or nested. In Section 3 we define a structure tree \mathcal{T} associated canonically with a nested set of separations of a graph G . In Section 4 we construct a tree-decomposition of G modelled

⁵Thus we shall prove that our counterexample is essentially the only one: all graphs not containing it have a unified tree-decomposition distinguishing all their blocks.

on \mathcal{T} , and study its parts. In Section 5 we find conditions under which, given a set \mathcal{S} of separations and a collection \mathcal{I} of \mathcal{S} -inseparable set of vertices, there is a nested subsystem of \mathcal{S} that still separates all the sets in \mathcal{I} . In Section 6, finally, we apply all this to the case of k -separations and k -blocks. We shall derive a central result, Theorem 6.3, which includes Theorems 1 and 2 as special cases.

2 Separations

Let $G = (V, E)$ be a finite graph. A *separation* of G is an ordered pair (A, B) such that $A, B \subseteq V$ and $G[A] \cup G[B] = G$. A separation (A, B) is *proper* if neither $A \setminus B$ nor $B \setminus A$ is empty. The *order* of a separation (A, B) is the cardinality of its *separator* $A \cap B$; the sets A, B are its *sides*. A separation of order k is a *k-separation*.

A separation (A, B) *separates* a set $I \subseteq V$ if I meets both $A \setminus B$ and $B \setminus A$. Two sets I_0, I_1 are *weakly separated* by a separation (A, B) if $I_i \subseteq A$ and $I_{1-i} \subseteq B$ for an $i \in \{0, 1\}$. They are *properly separated*, or simply *separated*, by (A, B) if in addition neither I_0 nor I_1 is contained in $A \cap B$.

Given a set \mathcal{S} of separations, we call a set of vertices *\mathcal{S} -inseparable* if no separation in \mathcal{S} separates it. A maximal \mathcal{S} -inseparable set of vertices is an *\mathcal{S} -block*, or simply a *block* if \mathcal{S} is fixed in the context.

Lemma 2.1. *Distinct \mathcal{S} -blocks b_1, b_2 are separated by some $(A, B) \in \mathcal{S}$.*

Proof. Since b_1 and b_2 are maximal \mathcal{S} -inseparable sets, $b := b_1 \cup b_2$ can be separated by some $(A, B) \in \mathcal{S}$. Then $b \setminus B \neq \emptyset \neq b \setminus A$, but being \mathcal{S} -inseparable, b_1 and b_2 are each contained in A or B . Hence (A, B) separates b_1 from b_2 . \square

A set of vertices is *small* with respect to \mathcal{S} if it is contained in the separator of some separation in \mathcal{S} . If \mathcal{S} is given from the context, we simply call such a set *small*. Note that if two sets are weakly but not properly separated by some separation in \mathcal{S} then at least one of them is small.

Let us look at how different separations of G can relate to each other. The set of all separations of G is partially ordered by

$$(A, B) \leq (C, D) \quad :\Leftrightarrow \quad A \subseteq C \text{ and } B \supseteq D. \quad (1)$$

Indeed, reflexivity, antisymmetry and transitivity follow easily from the corresponding properties of set inclusion on $\mathcal{P}(V)$. Note that changing the order in each pair reverses the relation:

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

Let (C, D) be any proper separation.

$$\begin{aligned} &\text{No proper separation } (A, B) \text{ is } \leq\text{-comparable with both } (C, D) \\ &\text{and } (D, C). \text{ In particular, } (C, D) \not\leq (D, C). \end{aligned} \quad (3)$$

Indeed, if $(A, B) \leq (C, D)$ and also $(A, B) \leq (D, C)$, then $A \subseteq C \subseteq B$ and hence $A \setminus B = \emptyset$, a contradiction. By (2), the other cases all reduce to this case by changing notation: just swap (A, B) with (B, A) or (C, D) or (D, C) .

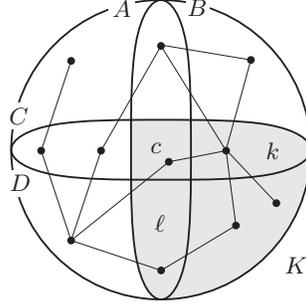


Figure 1: The cross-diagram $\{(A, B), (C, D)\}$ with centre c and a corner K and its links k, ℓ .

The way in which two separations relate to each other can be illustrated by a *cross-diagram* as in Figure 1. In view of such diagrams, we introduce the following terms for any set $\{(A, B), (C, D)\}$ of two separations, not necessarily distinct. The set $A \cap B \cap C \cap D$ is their *centre*, and $A \cap C, A \cap D, B \cap C, B \cap D$ are their *corners*. The corners $A \cap C$ and $B \cap D$ are *opposite*, as are the corners $A \cap D$ and $B \cap C$. Two corners that are not opposite are *adjacent*. The *link* between two adjacent corners is their intersection minus the centre. A corner minus its links and the centre is the *interior* of that corner; the rest – its two links and the centre – are its *boundary*. We shall write ∂K for the boundary of a corner K .

A corner forms a separation of G together with the union of the other three corners. We call these separations *corner separations*. For example, $(A \cap C, B \cup D)$ (in this order) is the corner separation for the corner $A \cap C$ in $\{(A, B), (C, D)\}$.

The four corner separations of a cross-diagram compare with the two separations forming it, and with the inverses of each other, in the obvious way:

$$\text{Any two separations } (A, B), (C, D) \text{ satisfy } (A \cap C, B \cup D) \leq (A, B). \quad (4)$$

$$\text{If } (I, J) \text{ and } (K, L) \text{ are distinct corner separations of the same cross-diagram, then } (I, J) \leq (L, K). \quad (5)$$

Inspection of the cross-diagram for (A, B) and (C, D) shows that $(A, B) \leq (C, D)$ if and only if the corner $A \cap D$ has an empty interior and empty links, i.e., the entire corner $A \cap D$ is contained in the centre:

$$(A, B) \leq (C, D) \Leftrightarrow A \cap D \subseteq B \cap C. \quad (6)$$

Another consequence of $(A, B) \leq (C, D)$ is that $A \cap B \subseteq C$ and $C \cap D \subseteq B$. So both separators live entirely on one side of the other separation.

A separation (A, B) is *tight* if every vertex of $A \cap B$ has a neighbour in $A \setminus B$ and another neighbour in $B \setminus A$. For tight separations, one can establish that $(A, B) \leq (C, D)$ by checking only one of the two inclusions in (1):

$$\text{If } (A, B) \text{ and } (C, D) \text{ are separations such that } A \subseteq C \text{ and } (C, D) \text{ is tight, then } (A, B) \leq (C, D). \quad (7)$$

Indeed, suppose $D \not\subseteq B$. Then as $A \subseteq C$, there is a vertex $x \in (C \cap D) \setminus B$. As (C, D) is tight, x has a neighbour $y \in D \setminus C$, but since $x \in A \setminus B$ we see that $y \in A$. So $A \setminus C \neq \emptyset$, contradicting our assumption.

Let us call (A, B) and (C, D) *nested*, and write $(A, B) \parallel (C, D)$, if (A, B) is comparable with (C, D) or with (D, C) under \leq . By (2), this is a symmetrical relation. For example, we saw in (4) and (5) that the corner separations of a cross-diagram are nested with the two separations forming it, as well as with each other.

Separations (A, B) and (C, D) that are not nested are said to *cross*; we then write $(A, B) \nparallel (C, D)$.

Nestedness is invariant under ‘flipping’ a separation: if $(A, B) \parallel (C, D)$ then also $(A, B) \parallel (D, C)$, by definition of \parallel , but also $(B, A) \parallel (C, D)$ by (2). Thus although nestedness is defined on the separations of G , we may think of it as a symmetrical relation on the unordered pairs $\{A, B\}$ such that (A, B) is a separation.

By (6), nested separations have a simple description in terms of cross-diagrams:

Two separations are nested if and only if one of their four corners has an empty interior and empty links. (8)

In particular:

Neither of two nested separations separates the separator of the other. (9)

The converse of (9) fails only if there is a corner with a non-empty interior whose links are both empty.

Although nestedness is reflexive and symmetric, it is not in general transitive. However when transitivity fails, we can still say something:

Lemma 2.2. *If $(A, B) \parallel (C, D)$ and $(C, D) \parallel (E, F)$ but $(A, B) \nparallel (E, F)$, then (C, D) is nested with every corner separation of $\{(A, B), (E, F)\}$, and for one corner separation (I, J) we have either $(C, D) \leq (I, J)$ or $(D, C) \leq (I, J)$.*

Proof. Changing notation as necessary, we may assume that $(A, B) \leq (C, D)$, and that (C, D) is comparable with (E, F) .⁶ If $(C, D) \leq (E, F)$ we have $(A, B) \leq (E, F)$, contrary to our assumption. Hence $(C, D) \geq (E, F)$, or equivalently by (2), $(D, C) \leq (F, E)$. As also $(D, C) \leq (B, A)$, we thus have $D \subseteq F \cap B$ and $C \supseteq E \cup A$ and therefore

$$(D, C) \leq (F \cap B, E \cup A) \underset{(5)}{\leq} (L, K)$$

for each of the other three corner separations (K, L) of $\{(A, B), (E, F)\}$. \square

⁶Note that such change of notation will not affect the set of corner separations of the cross-diagram of (A, B) and (E, F) , nor the nestedness (or not) of (C, D) with those corner separations.

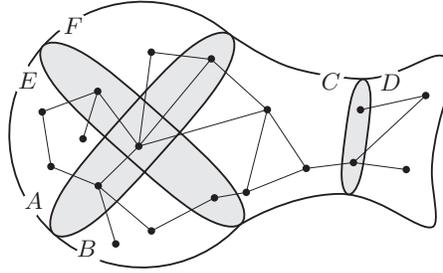


Figure 2: Separations as in Lemma 2.2

Figure 2 shows an example of three separations witnessing the non-transitivity of nestedness. Its main purpose, however, is to illustrate the use of Lemma 2.2. We shall often be considering which of two crossing separations, such as (A, B) and (E, F) in the example, we should adopt for a desired collection of nested separations already containing some separations such as (C, D) . The lemma then tells us that we can opt to take neither, but instead choose a suitable corner separation.

Note that there are two ways in which three separations can be pairwise nested. One is that they or their inverses form a chain under \leq . But there is also another way, which will be important later; this is illustrated in Figure 3.

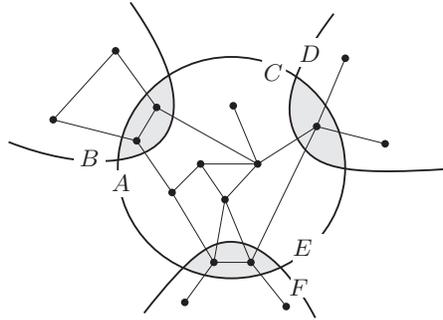


Figure 3: Three nested separations not coming from a \leq -chain

We need one more lemma.

Lemma 2.3. *Let \mathcal{N} be a set of separations of G that are pairwise nested. Let (A, B) and (C, D) be two further separations, each nested with all the separations in \mathcal{N} . Assume that (A, B) separates an \mathcal{N} -block b , and that (C, D) separates an \mathcal{N} -block $b' \neq b$. Then $(A, B) \parallel (C, D)$. Moreover, $A \cap B \subseteq b$ and $C \cap D \subseteq b'$.*

Proof. By Lemma 2.1, there is a separation $(E, F) \in \mathcal{N}$ with $b \subseteq E$ and $b' \subseteq F$. Suppose $(A, B) \not\parallel (C, D)$. By symmetry and Lemma 2.2 we may assume that

$$(E, F) \leq (A \cap C, B \cup D).$$

But then $b \subseteq E \subseteq A \cap C \subseteq A$, contradicting the fact that (A, B) separates b . Hence $(A, B) \parallel (C, D)$, as claimed.

If $A \cap B \not\subseteq b$, then there is a $(K, L) \in \mathcal{N}$ which separates $b \cup (A \cap B)$. We may assume that $b \subseteq L$ and that $A \cap B \not\subseteq L$. The latter implies that $(K, L) \not\leq (A, B)$ and $(K, L) \not\leq (B, A)$. So $(K, L) \parallel (A, B)$ implies that either $(L, K) \leq (A, B)$ or $(L, K) \leq (B, A)$. Thus $b \subseteq L \subseteq A$ or $b \subseteq L \subseteq B$, a contradiction to the fact that (A, B) separates b . Similarly we obtain $C \cap D \subseteq b'$. \square

3 Nested separation systems and tree structure

A set \mathcal{S} of separations is *symmetric* if $(A, B) \in \mathcal{S}$ implies $(B, A) \in \mathcal{S}$, and *nested* if every two separations in \mathcal{S} are nested. Any symmetric set of proper separations is a *separation system*. Throughout this section and the next, we consider a fixed nested separation system \mathcal{N} of our graph G .

Our aim in this section will be to describe \mathcal{N} by way of a *structure tree* $\mathcal{T} = \mathcal{T}(\mathcal{N})$, whose edges will correspond to the separations in \mathcal{N} . Its nodes⁷ will correspond to subgraphs of G . Every automorphism of G that leaves \mathcal{N} invariant will also act on \mathcal{T} . Although our notion of a separation system differs from that of Dunwoody and Krön [6, 5], the main ideas of how to describe a nested system by a structure tree can already be found there.

Our main task in the construction of \mathcal{T} will be to define its nodes. They will be the equivalence classes of the following equivalence relation \sim on \mathcal{N} , induced by the ordering \leq from (1):

$$(A, B) \sim (C, D) \quad :\Leftrightarrow \quad \begin{cases} (A, B) = (C, D) \text{ or} \\ (B, A) \text{ is a predecessor of } (C, D) \text{ in } (\mathcal{N}, \leq). \end{cases} \quad (10)$$

(Recall that, in a partial order (P, \leq) , an element $x \in P$ is a *predecessor* of an element $z \in P$ if $x < z$ but there is no $y \in P$ with $x < y < z$.)

Before we prove that this is indeed an equivalence relation, it may help to look at an example: the set of vertices in the centre of Figure 3 will be the node of \mathcal{T} represented by each of the equivalent nested separations (A, B) , (C, D) and (E, F) .

Lemma 3.1. *The relation \sim is an equivalence relation on \mathcal{N} .*

Proof. Reflexivity holds by definition, and symmetry follows from (2). To show transitivity assume that $(A, B) \sim (C, D)$ and $(C, D) \sim (E, F)$, and that all these separations are distinct. Thus,

- (i) (B, A) is a predecessor of (C, D) ;
- (ii) (D, C) is a predecessor of (E, F) .

And by (2) also

- (iii) (D, C) is a predecessor of (A, B) ;
- (iv) (F, E) is a predecessor of (C, D) .

⁷While our graphs G have *vertices*, structure trees will have *nodes*.

By (ii) and (iii), (A, B) is incomparable with (E, F) . Hence, since \mathcal{N} is nested, (B, A) is comparable with (E, F) . If $(E, F) \leq (B, A)$ then by (i) and (ii), $(D, C) \leq (C, D)$, which contradicts (3) (recall that all separations in a separation system are required to be proper). Thus $(B, A) < (E, F)$, as desired.

Suppose there is a separation $(X, Y) \in \mathcal{N}$ with $(B, A) < (X, Y) < (E, F)$. As \mathcal{N} is nested, (X, Y) is comparable with either (C, D) or (D, C) . By (i) and (ii), $(X, Y) \not\leq (C, D)$ and $(D, C) \not\leq (X, Y)$. Now if $(C, D) \leq (X, Y) < (E, F)$ then by (iv), (C, D) is comparable to both (E, F) and (F, E) , contradicting (3). Finally, if $(D, C) \geq (X, Y) > (B, A)$, then by (iii), (D, C) is comparable to both (B, A) and (A, B) , again contradicting (3). We have thus shown that (B, A) is a predecessor of (E, F) , implying that $(A, B) \sim (E, F)$ as claimed. \square

Note that, by (3), the definition of equivalence implies:

$$\text{Distinct equivalent proper separations are incomparable under } \leq. \quad (11)$$

We can now define the nodes of $\mathcal{T} = \mathcal{T}(\mathcal{N})$ as planned, as the equivalence classes of \sim :

$$V(\mathcal{T}) := \{[(A, B)] : (A, B) \in \mathcal{N}\}.$$

Having defined the nodes of \mathcal{T} , let us define its edges. For every separation $(A, B) \in \mathcal{N}$ we shall have one edge, joining the nodes represented by (A, B) and (B, A) , respectively. To facilitate notation later, we formally give \mathcal{T} the abstract edge set

$$E(\mathcal{T}) := \{\{(A, B), (B, A)\} \mid (A, B) \in \mathcal{N}\}$$

and declare an edge e to be incident with a node $t \in V(\mathcal{T})$ whenever $e \cap t \neq \emptyset$ (so that the edge $\{(A, B), (B, A)\}$ of \mathcal{T} joins its nodes $[(A, B)]$ and $[(B, A)]$). We have thus, so far, defined a multigraph \mathcal{T} .

As $(A, B) \not\sim (B, A)$ by definition of \sim , our multigraph \mathcal{T} has no loops. Whenever an edge e is incident with a node t , the non-empty set $e \cap t$ that witnesses this is a singleton set containing one separation. We denote this separation by $(e \cap t)$. Every separation $(A, B) \in \mathcal{N}$ occurs as such an $(e \cap t)$, with $t = [(A, B)]$ and $e = \{(A, B), (B, A)\}$. Thus,

$$\text{Every node } t \text{ of } \mathcal{T} \text{ is the set of all the separations } (e \cap t) \text{ such that } e \text{ is incident with } t. \text{ In particular, } t \text{ has degree } |t| \text{ in } \mathcal{T}. \quad (12)$$

Our next aim is to show that \mathcal{T} is a tree.

Lemma 3.2. *Let $W = t_1 e_1 t_2 e_2 t_3$ be a walk in \mathcal{T} with $e_1 \neq e_2$. Then $(e_1 \cap t_1)$ is a predecessor of $(e_2 \cap t_2)$.*

Proof. Let $(e_1 \cap t_1) = (A, B)$ and $(e_2 \cap t_2) = (C, D)$. Then $(B, A) = (e_1 \cap t_2)$ and $(B, A) \sim (C, D)$. Since $e_1 \neq e_2$ we have $(B, A) \neq (C, D)$. Thus, (A, B) is a predecessor of (C, D) by definition of \sim . \square

And conversely:

Lemma 3.3. *Let $(E_0, F_0), \dots, (E_k, F_k)$ be separations in \mathcal{N} such that each (E_{i-1}, F_{i-1}) is a predecessor of (E_i, F_i) in (\mathcal{N}, \leq) . Then $[(E_0, F_0)], \dots, [(E_k, F_k)]$ are the nodes of a walk in \mathcal{T} , in this order.*

Proof. By definition of \sim , we know that $(F_{i-1}, E_{i-1}) \sim (E_i, F_i)$. Hence for all $i = 1, \dots, k$, the edge $\{(E_{i-1}, F_{i-1}), (F_{i-1}, E_{i-1})\}$ of \mathcal{T} joins the node $[(E_{i-1}, F_{i-1})]$ to the node $[(E_i, F_i)] = [(F_{i-1}, E_{i-1})]$. \square

Theorem 3.4. *The multigraph $\mathcal{T}(\mathcal{N})$ is a tree.*

Proof. We have seen that \mathcal{T} is loopless. Suppose that \mathcal{T} contains a cycle $t_1 e_1 \cdots t_{k-1} e_{k-1} t_k$, with $t_1 = t_k$ and $k > 2$. Applying Lemma 3.2 $(k-1)$ times yields

$$(A, B) := (e_1 \cap t_1) < \dots < (e_{k-1} \cap t_{k-1}) < (e_1 \cap t_k) = (A, B),$$

a contradiction. Thus, \mathcal{T} is acyclic; in particular, it has no parallel edges.

It remains to show that \mathcal{T} contains a path between any two given nodes $[(A, B)]$ and $[(C, D)]$. As \mathcal{N} is nested, we know that (A, B) is comparable with either (C, D) or (D, C) . Since $[(C, D)]$ and $[(D, C)]$ are adjacent, it suffices to construct a walk between $[(A, B)]$ and one of them. Swapping the names for C and D if necessary, we may thus assume that (A, B) is comparable with (C, D) . Reversing the direction of our walk if necessary, we may further assume that $(A, B) < (C, D)$. Since our graph G is finite, there is a chain

$$(A, B) = (E_0, F_0) < \dots < (E_k, F_k) = (C, D)$$

such that (E_{i-1}, F_{i-1}) is a predecessor of (E_i, F_i) , for every $i = 1, \dots, k$. By Lemma 3.3, \mathcal{T} contains the desired path from $[(A, B)]$ to $[(C, D)]$. \square

Corollary 3.5. *If \mathcal{N} is invariant under a group $\Gamma \leq \text{Aut}(G)$ of automorphisms of G , then Γ also acts on \mathcal{T} as a group of automorphisms.*

Proof. Any automorphism α of G maps separations to separations, and preserves their partial ordering defined in (1). If both α and α^{-1} map separations from \mathcal{N} to separations in \mathcal{N} , then α also preserves the equivalence of separations under \sim . Hence Γ , as stated, acts on the nodes of \mathcal{T} and preserves their adjacencies and non-adjacencies. \square

4 From structure trees to tree-decompositions

Throughout this section, \mathcal{N} continues to be an arbitrary nested separation system of our graph G . Our aim now is to show that G has a tree-decomposition, in the sense of Robertson and Seymour, with the structure tree $\mathcal{T} = \mathcal{T}(\mathcal{N})$ defined in Section 3 as its decomposition tree. The separations of G associated with the edges of this decomposition tree⁸ will be precisely the separations in \mathcal{N} identified by those edges in the original definition of \mathcal{T} .

⁸as in the theory of tree-decompositions, see e.g. [3, Lemma 12.3.1]

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 .

To define our desired tree-decomposition $(\mathcal{T}, \mathcal{V})$, we thus have to define the family $\mathcal{V} = (V_t)_{t \in \mathcal{T}}$ of its parts: with every node t of \mathcal{T} we have to associate a set V_t of vertices of G . We define these as follows:

$$V_t := \bigcap \{ A \mid (A, B) \in t \} \quad (13)$$

Example 1. Assume that G is connected, and consider as \mathcal{N} the nested set of all proper 1-separations (A, B) and (B, A) such that $A \setminus B$ is connected in G . Then \mathcal{T} is very similar to the block-cutvertex tree of G : its nodes will be the blocks in the usual sense (maximal 2-connected subgraphs or bridges) plus those cutvertices that lie in at least three blocks.

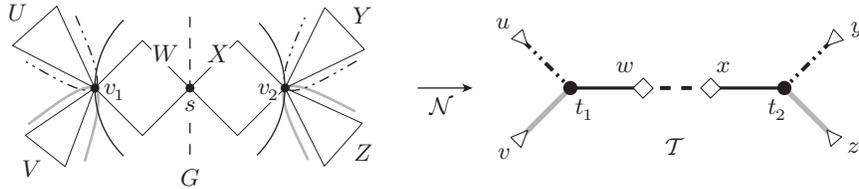


Figure 4: \mathcal{T} has an edge for every separation in \mathcal{N} . Its nodes correspond to the blocks and some of the cutvertices of G .

In Figure 4, this separation system \mathcal{N} contains all the proper 1-separations of G . The separation (A, B) defined by the cutvertex s , with $A := U \cup V \cup W$ and $B := X \cup Y \cup Z$ say, defines the edge $\{(A, B), (B, A)\}$ of \mathcal{T} joining its nodes $w = [(A, B)]$ and $x = [(B, A)]$.

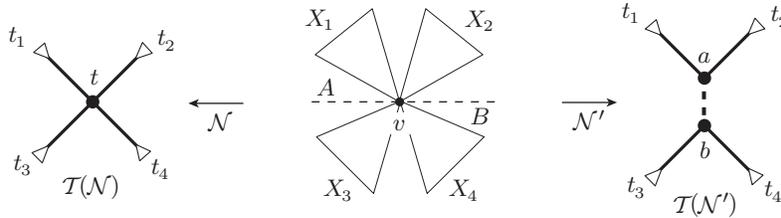


Figure 5: $\mathcal{T}' = \mathcal{T}(\mathcal{N}')$ has distinct nodes a, b whose parts in the tree-decomposition $(\mathcal{T}', \mathcal{V})$ coincide: $V_a = \{v\} = V_b$.

In Figure 5 we can add to \mathcal{N} one of the two crossing 1-separations not in \mathcal{N} (together with its inverse), to obtain a set \mathcal{N}' of separations that is still nested.

For example, let

$$\mathcal{N}' := \mathcal{N} \cup \{(A, B), (B, A)\}$$

with $A := X_1 \cup X_2$ and $B := X_3 \cup X_4$. This causes the central node t of \mathcal{T} to split into two nodes $a = [(A, B)]$ and $b = [(B, A)]$ joined by the new edge $\{(A, B), (B, A)\}$. However the new nodes a, b still define the same part of the tree-decomposition of G as t did before: $V_a = V_b = V_t = \{v\}$.

Before we prove that $(\mathcal{T}, \mathcal{V})$ is indeed a tree-decomposition, let us collect some information about its parts V_t , the vertex sets defined in (13).

Lemma 4.1. *Every V_t is \mathcal{N} -inseparable.*

Proof. Let us show that a given separation $(C, D) \in \mathcal{N}$ does not separate V_t . Pick $(A, B) \in t$. Since \mathcal{N} is nested, and swapping the names of C and D if necessary, we may assume that (A, B) is \leq -comparable with (C, D) . If $(A, B) \leq (C, D)$ then $V_t \subseteq A \subseteq C$, so (C, D) does not separate V_t . If $(C, D) < (A, B)$, there is a \leq -predecessor (E, F) of (A, B) with $(C, D) \leq (E, F)$. Then $(F, E) \sim (A, B)$ and hence $V_t \subseteq F \subseteq D$, so again (C, D) does not separate V_t . \square

The sets V_t will come in two types: they can be

- \mathcal{N} -blocks (that is, maximal \mathcal{N} -inseparable sets of vertices), or
- ‘hubs’ (defined below).

Nodes $t \in \mathcal{T}$ such that V_t is an \mathcal{N} -block are *block nodes*. A node $t \in \mathcal{T}$ such that $V_t = A \cap B$ for some $(A, B) \in t$ is a *hub node* (and V_t a *hub*).

In Example 1, the \mathcal{N} -blocks were the (usual) blocks of G ; the hubs were singleton sets consisting of a cutvertex. Example 2 will show that t can be a hub node and a block node at the same time. Every hub is a subset of a block: by (9), hubs are \mathcal{N} -inseparable, so they extend to maximal \mathcal{N} -inseparable sets.

Hubs can contain each other properly (Example 2 below). But a hub V_t cannot be properly contained in a separator $A \cap B$ of any $(A, B) \in t$. Let us prove this without assuming that V_t is a hub:

Lemma 4.2. *Whenever $(A, B) \in t \in \mathcal{T}$, we have $A \cap B \subseteq V_t$. In particular, if $V_t \subseteq A \cap B$, then $V_t = A \cap B$ is a hub with hub node t .*

Proof. Consider any vertex $v \in (A \cap B) \setminus V_t$. By definition of V_t , there exists a separation $(C, D) \in t$ such that $v \notin C$. This contradicts the fact that $B \subseteq C$ since $(A, B) \sim (C, D)$. \square

Lemma 4.3. *Every node of \mathcal{T} is either a block node or a hub node.*

Proof. Suppose $t \in \mathcal{T}$ is not a hub node; we show that t is a block node. By Lemma 4.1, V_t is \mathcal{N} -inseparable. We show that V_t is maximal in $V(G)$ with this property: that for every vertex $x \notin V_t$ the set $V_t \cup \{x\}$ is not \mathcal{N} -inseparable.

By definition of V_t , any vertex $x \notin V_t$ lies in $B \setminus A$ for some $(A, B) \in t$. Since t is not a hub node, Lemma 4.2 implies that $V_t \not\subseteq A \cap B$. As $V_t \subseteq A$, this means that V_t has a vertex in $A \setminus B$. Hence (A, B) separates $V_t \cup \{x\}$, as desired. \square

Conversely, all the \mathcal{N} -blocks of G will be parts of our tree-decomposition:

Lemma 4.4. *Every \mathcal{N} -block is the set V_t for a node t of \mathcal{T} .*

Proof. Consider an arbitrary \mathcal{N} -block b .

Suppose first that b is small. Then there exists a separation $(A, B) \in \mathcal{N}$ with $b \subseteq A \cap B$. As \mathcal{N} is nested, $A \cap B$ is \mathcal{N} -inseparable by (9), so in fact $b = A \cap B$ by the maximality of b . We show that $b = V_t$ for $t = [(A, B)]$. By Lemma 4.2, it suffices to show that $V_t \subseteq b = A \cap B$. As $V_t \subseteq A$ by definition of V_t , we only need to show that $V_t \subseteq B$. Suppose there is an $x \in V_t \setminus B$. As $x \notin A \cap B = b$, the maximality of b implies that there exists a separation $(E, F) \in \mathcal{N}$ such that

$$F \not\supseteq b \subseteq E \text{ and } x \in F \setminus E \quad (*)$$

(compare the proof of Lemma 2.1). By (*), all corners of the cross-diagram $\{(A, B), (E, F)\}$ other than $B \cap F$ contain vertices not in the centre. Hence by (8), the only way in which (A, B) and (E, F) can be nested is that $B \cap F$ does lie in the centre, i.e. that $(B, A) \leq (E, F)$. Since $(B, A) \neq (E, F)$, by (*) and $b = A \cap B$, this means that (B, A) has a successor $(C, D) \leq (E, F)$. But then $(C, D) \sim (A, B)$ and $x \notin E \supseteq C \supseteq V_t$, a contradiction.

Suppose now that b is not small. We shall prove that $b = V_t$ for $t = t(b)$, where $t(b)$ is defined as the set of separations (A, B) that are minimal with $b \subseteq A$. Let us show first that $t(b)$ is indeed an equivalence class, i.e., that the separations in $t(b)$ are equivalent to each other but not to any other separation in \mathcal{N} .

Given distinct $(A, B), (C, D) \in t(b)$, let us show that $(A, B) \sim (C, D)$. Since both (A, B) and (C, D) are minimal as in the definition of $t(b)$, they are incomparable. But as elements of \mathcal{N} they are nested, so (A, B) is comparable with (D, C) . If $(A, B) \leq (D, C)$ then $b \subseteq A \cap C \subseteq D \cap C$, which contradicts our assumption that b is not small. Hence $(D, C) < (A, B)$. To show that (D, C) is a predecessor of (A, B) , suppose there exists a separation $(E, F) \in \mathcal{N}$ such that $(D, C) < (E, F) < (A, B)$. This contradicts the minimality either of (A, B) , if $b \subseteq E$, or of (C, D) , if $b \subseteq F$. Thus, $(C, D) \sim (A, B)$ as desired.

Conversely, we have to show that every $(E, F) \in \mathcal{N}$ equivalent to some $(A, B) \in t(b)$ also lies in $t(b)$. As $(E, F) \sim (A, B)$, we may assume that $(F, E) < (A, B)$. Then $b \not\subseteq F$ by the minimality of (A, B) as an element of $t(b)$, so $b \subseteq E$. To show that (E, F) is minimal with this property, suppose that $b \subseteq X$ also for some $(X, Y) \in \mathcal{N}$ with $(X, Y) < (E, F)$. Then (X, Y) is incomparable with (A, B) : by (11) we cannot have $(A, B) \leq (X, Y) < (E, F)$, and we cannot have $(X, Y) < (A, B)$ by the minimality of (A, B) as an element of $t(b)$. But $(X, Y) \parallel (A, B)$, so (X, Y) must be comparable with (B, A) . Yet if $(X, Y) \leq (B, A)$, then $b \subseteq X \cap A \subseteq B \cap A$, contradicting our assumption that

b is not small, while $(B, A) < (X, Y) < (E, F)$ is impossible, since (B, A) is a predecessor of (E, F) .

Hence $t(b)$ is indeed an equivalence class, i.e., $t(b) \in V(\mathcal{T})$. By definition of $t(b)$, we have $b \subseteq \bigcap \{A \mid (A, B) \in t(b)\} = V_{t(b)}$. The converse inclusion follows from the maximality of b as an \mathcal{N} -inseparable set. \square

We have seen so far that the parts V_t of our intended tree-decomposition associated with \mathcal{N} are all the \mathcal{N} -blocks of G , plus some hubs. The following proposition shows what has earned them their name:

Proposition 4.5. *A hub node t has degree at least 3 in \mathcal{T} , unless it has the form $t = \{(A, B), (C, D)\}$ with $A \supseteq D$ and $B = C$ (in which case it has degree 2).*

Proof. Let $(A, B) \in t$ be such that $V_t = A \cap B$. As $(A, B) \in t$ but $V_t \neq A$, we have $d(t) = |t| \geq 2$; cf. (12). Suppose that $d(t) = 2$, say $t = \{(A, B), (C, D)\}$. Then $B \subseteq C$ by definition of \sim , and $C \setminus B = (C \cap A) \setminus B = V_t \setminus B = \emptyset$ by definition of V_t and $V_t \subseteq A \cap B$. So $B = C$. As (A, B) and (C, D) are equivalent but not equal, this implies $D \subsetneq A$. \square

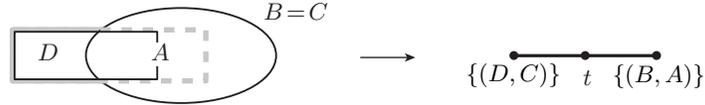


Figure 6: A hub node $t = \{(A, B), (C, D)\}$ of degree 2

Figure 6 shows that the exceptional situation from Proposition 4.5 can indeed occur. In the example, we have $\mathcal{N} = \{(A, B), (B, A), (C, D), (D, C)\}$ with $B = C$ and $D \subsetneq A$. The structure tree \mathcal{T} is a path between two block nodes $\{(D, C)\}$ and $\{(B, A)\}$ with a central hub node $t = \{(A, B), (C, D)\}$, whose set $V_t = A \cap B$ is not a block since it is properly contained in the \mathcal{N} -inseparable set $B = C$.

Our last example answers some further questions about the possible relationships between blocks and hubs that will naturally come to mind:

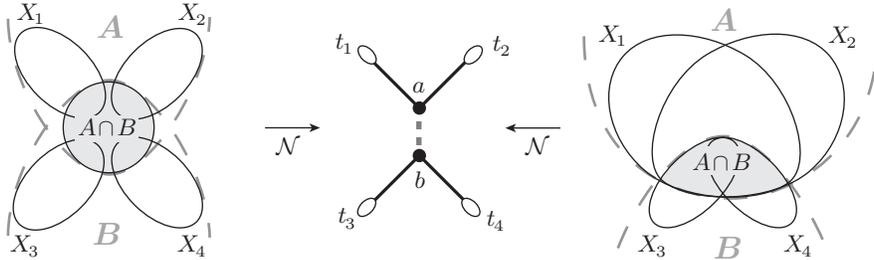


Figure 7: The two nested separation systems of Example 2, and their common structure tree

Example 2. Consider the vertex sets X_1, \dots, X_4 shown on the left of Figure 7. Let A be a superset of $X_1 \cup X_2$ and B a superset of $X_3 \cup X_4$, so that $A \cap B \not\subseteq X_1 \cup \dots \cup X_4$ and different X_i do not meet outside $A \cap B$. Let \mathcal{N} consist of (A, B) , (B, A) , and $(X_1, Y_1), \dots, (X_4, Y_4)$ and their inverses (Y_i, X_i) , where $Y_i := (A \cap B) \cup \bigcup_{j \neq i} X_j$. The structure tree $\mathcal{T} = \mathcal{T}(\mathcal{N})$ has four block nodes t_1, \dots, t_4 , with $t_i = [(X_i, Y_i)]$ and $V_{t_i} = X_i$, and two central hub nodes

$$a = \{(A, B), (Y_1, X_1), (Y_2, X_2)\} \quad \text{and} \quad b = \{(B, A), (Y_3, X_3), (Y_4, X_4)\}$$

joined by the edge $\{(A, B), (B, A)\}$. The hubs corresponding to a and b coincide: they are $V_a = A \cap B = V_b$, which is also a block.

Let us now modify this example by enlarging X_1 and X_2 so that they meet outside $A \cap B$ and each contain $A \cap B$. Thus, $A = X_1 \cup X_2$. Let us also shrink B a little, down to $B = X_3 \cup X_4$ (Fig. 7, right). The structure tree \mathcal{T} remains unchanged by these modifications, but the corresponding sets V_t have changed:

$$V_b = A \cap B \subsetneq X_1 \cap X_2 = X_1 \cap Y_1 = X_2 \cap Y_2 = V_a,$$

and neither of them is a block, because both are properly contained in X_1 , which is also \mathcal{N} -inseparable.

Our next lemma shows that deleting a separation from our nested system \mathcal{N} corresponds to contracting an edge in the structure tree $\mathcal{T}(\mathcal{N})$. For a separation (A, B) that belongs to different systems, we write $[(A, B)]_{\mathcal{N}}$ to indicate in which system \mathcal{N} we are taking the equivalence class.

Lemma 4.6. *Given $(A, B) \in \mathcal{N}$, the tree $\mathcal{T}' := \mathcal{T}(\mathcal{N}')$ for*

$$\mathcal{N}' = \mathcal{N} \setminus \{(A, B), (B, A)\}$$

arises from $\mathcal{T} = \mathcal{T}(\mathcal{N})$ by contracting the edge $e = \{(A, B), (B, A)\}$. The contracted node z of \mathcal{T}' satisfies $z = x \cup y \setminus e$ and $V_z = V_x \cup V_y$, where $x = [(A, B)]_{\mathcal{N}}$ and $y = [(B, A)]_{\mathcal{N}}$, and $V(\mathcal{T}') \setminus \{z\} = V(\mathcal{T}) \setminus \{x, y\}$.⁹

Proof. To see that $V(\mathcal{T}') \setminus \{z\} = V(\mathcal{T}) \setminus \{x, y\}$ and $z = x \cup y \setminus e$, we have to show for all $(C, D) \in \mathcal{N}'$ that $[(C, D)]_{\mathcal{N}'} = [(C, D)]_{\mathcal{N}}$ unless $[(C, D)]_{\mathcal{N}} \in \{x, y\}$, in which case $[(C, D)]_{\mathcal{N}'} = x \cup y \setminus e$. In other words, we have to show:

$$\text{Two separations } (C, D), (E, F) \in \mathcal{N}' \text{ are equivalent in } \mathcal{N}' \text{ if and only if they are equivalent in } \mathcal{N} \text{ or are both in } x \cup y \setminus e. \quad (*)$$

Our further claim that $\mathcal{T}' = \mathcal{T}/e$, i.e. that the node-edge incidences in \mathcal{T}' arise from those in \mathcal{T} as defined for graph minors, will follow immediately from the definition of these incidences in \mathcal{T} and \mathcal{T}' .

Let us prove the backward implication of $(*)$ first. As $\mathcal{N}' \subseteq \mathcal{N}$, predecessors in (\mathcal{N}, \leq) are still predecessors in \mathcal{N}' , and hence $(C, D) \sim_{\mathcal{N}'} (E, F)$

⁹The last identity says more than that there exists a canonical bijection between $V(\mathcal{T}') \setminus \{z\}$ and $V(\mathcal{T}) \setminus \{x, y\}$: it says that the nodes of $\mathcal{T} - \{x, y\}$ and $\mathcal{T}' - z$ are the same also as sets of separations.

implies $(C, D) \sim_{\mathcal{N}'} (E, F)$. Moreover if $(C, D) \in x$ and $(E, F) \in y$ then, in \mathcal{N} , (D, C) is a predecessor of (A, B) and (A, B) is a predecessor of (E, F) . In \mathcal{N}' , then, (D, C) is a predecessor of (E, F) , since by Lemma 3.3 and Theorem 3.4 there is no separation $(A', B') \neq (A, B)$ in \mathcal{N} that is both a successor of (D, C) and a predecessor of (E, F) . Hence $(C, D) \sim_{\mathcal{N}'} (E, F)$.

For the forward implication in $(*)$ note that if (D, C) is a predecessor of (E, F) in \mathcal{N}' but not in \mathcal{N} , then in \mathcal{N} we have a sequence of predecessors $(D, C) < (A, B) < (E, F)$ or $(D, C) < (B, A) < (E, F)$. Then one of (C, D) and (E, F) lies in x and the other in y , as desired.

It remains to show that $V_z = V_x \cup V_y$. Consider the sets

$$x' := x \setminus \{(A, B)\} \quad \text{and} \quad y' := y \setminus \{(B, A)\};$$

then $z = y' \cup x'$. Since all $(E, F) \in x'$ are equivalent to (A, B) but not equal to it, we have $(B, A) \leq (E, F)$ for all those separations. That is,

$$B \subseteq \bigcap_{(E,F) \in x'} E = V_{x'}. \quad (14)$$

By definition of V_x we have $V_x = V_{x'} \cap A$. Hence (14) yields $V_{x'} = V_x \cup (B \setminus A)$, and since $A \cap B \subseteq V_x$ by Lemma 4.2, we have $V_{x'} = V_x \cup B$. An analogous argument yields

$$V_{y'} = \bigcap_{(E,F) \in y'} E = V_y \cup A.$$

Hence,

$$\begin{aligned} V_z &= \bigcap_{(E,F) \in z} E \\ &= V_{x'} \cap V_{y'} \\ &= (V_x \cup B) \cap (V_y \cup A) \\ &= (V_x \cap V_y) \cup (V_x \cap A) \cup (V_y \cap B) \cup (B \cap A) \\ &= (V_x \cap V_y) \cup V_x \cup V_y \cup (B \cap A) \\ &= V_x \cup V_y. \end{aligned} \quad \square$$

Every edge e of \mathcal{T} separates \mathcal{T} into two components. The vertex sets V_t for the nodes t in these components induce a corresponding separation of G , as in [3, Lemma 12.3.1]. This is the separation that defined e :

Lemma 4.7. *Given any separation $(A, B) \in \mathcal{N}$, consider the corresponding edge $e = \{(A, B), (B, A)\}$ of $\mathcal{T} = \mathcal{T}(\mathcal{N})$. Let \mathcal{T}_A denote the component of $\mathcal{T} - e$ that contains the node $[(A, B)]$, and let \mathcal{T}_B be the other component. Then $\bigcup_{t \in \mathcal{T}_A} V_t = A$ and $\bigcup_{t \in \mathcal{T}_B} V_t = B$.*

Proof. We apply induction on $|E(\mathcal{T})|$. If \mathcal{T} consists of a single edge, the assertion is immediate from the definition of \mathcal{T} . Assume now that $|E(\mathcal{T})| > 1$. In particular, there is an edge $e^* = xy \neq e$.

Consider $\mathcal{N}' := \mathcal{N} \setminus e^*$, and let $\mathcal{T}' := \mathcal{T}(\mathcal{N}')$. Then $\mathcal{T}' = \mathcal{T}/e^*$, by Lemma 4.6. Let z be the node of \mathcal{T}' contracted from e^* . Define \mathcal{T}'_A as the component of $\mathcal{T}' - e$ that contains the node $[(A, B)]$, and let \mathcal{T}'_B be the other component. We may assume $e^* \in \mathcal{T}_A$. Then

$$V(\mathcal{T}_A) \setminus \{x, y\} = V(\mathcal{T}'_A) \setminus \{z\} \text{ and } V(\mathcal{T}_B) = V(\mathcal{T}'_B).$$

As $V_z = V_x \cup V_y$ by Lemma 4.6, we can use the induction hypothesis to deduce that

$$\bigcup_{t \in \mathcal{T}_A} V_t = \bigcup_{t \in \mathcal{T}'_A} V_t = A \quad \text{and} \quad \bigcup_{t \in \mathcal{T}_B} V_t = \bigcup_{t \in \mathcal{T}'_B} V_t = B,$$

as claimed. \square

Let us summarize some of our findings from this section. Recall that \mathcal{N} is an arbitrary nested separation system of an arbitrary finite graph G . Let $\mathcal{T} := \mathcal{T}(\mathcal{N})$ be the structure tree associated with \mathcal{N} as in Section 3, and let $\mathcal{V} := (V_t)_{t \in \mathcal{T}}$ be defined by (13). Let us call the separations of G that correspond as in [3, Lemma 12.3.1] to the edges of the decomposition tree of a tree-decomposition of G the separations *induced by* this tree-decomposition.

Theorem 4.8. *The pair $(\mathcal{T}, \mathcal{V})$ is a tree-decomposition of G .*

- (i) *Every \mathcal{N} -block is a part of the decomposition.*
- (ii) *Every part of the decomposition is either an \mathcal{N} -block or a hub.*
- (iii) *The separations of G induced by the decomposition are precisely those in \mathcal{N} .*
- (iv) *Every $\mathcal{N}' \subseteq \mathcal{N}$ satisfies $(\mathcal{T}', \mathcal{V}') \preceq (\mathcal{T}, \mathcal{V})$ for $\mathcal{T}' = \mathcal{T}(\mathcal{N}')$ and $\mathcal{V}' = V(\mathcal{T}')$.¹⁰*

Proof. Of the three axioms for a tree-decomposition, (T1) and (T2) follow from Lemma 4.4, because single vertices and edges form \mathcal{N} -inseparable vertex sets, which extend to \mathcal{N} -blocks. For the proof of (T3), let $e = \{(A, B), (B, A)\}$ be an edge at t_2 on the t_1 - t_3 path in \mathcal{T} . Since e separates t_1 from t_3 in \mathcal{T} , Lemmas 4.7 and 4.2 imply that $V_{t_1} \cap V_{t_3} \subseteq A \cap B \subseteq V_{t_2}$.

Statement (i) is Lemma 4.4. Assertion (ii) is Lemma 4.3. Assertion (iii) follows from Lemma 4.7 and the definition of the edges of \mathcal{T} . Statement (iv) follows by repeated application of Lemma 4.6. \square

5 Extracting nested separation systems

Our aim in this section will be to find inside a given separation system \mathcal{S} a nested subsystem \mathcal{N} that can still distinguish the elements of some given set \mathcal{I} of \mathcal{S} -inseparable sets of vertices. As we saw in Sections 3 and 4, such a nested subsystem will then define a tree-decomposition of G , and the sets from \mathcal{I} will come to lie in different parts of that decomposition.

¹⁰See the Introduction for the definition of $(\mathcal{T}', \mathcal{V}') \preceq (\mathcal{T}, \mathcal{V})$.

This cannot be done for all choices of \mathcal{S} and \mathcal{I} . Indeed, consider the following example of where such a nested subsystem does not exist. Let G be the 3×3 -grid, let \mathcal{S} consist of the two 3-separations cutting along the horizontal and the vertical symmetry axis, and let \mathcal{I} consist of the four corners of the resulting cross-diagram. Each of these is \mathcal{S} -inseparable, and any two of them can be separated by a separation in \mathcal{S} . But since the two separations in \mathcal{S} cross, any nested subsystem contains at most one of them, and thus fails to separate some sets from \mathcal{I} .

However, we shall prove that the desired nested subsystem does exist if \mathcal{S} and \mathcal{I} satisfy the following condition. Given a separation system \mathcal{S} and a set \mathcal{I} of \mathcal{S} -inseparable sets, let us say that \mathcal{S} *separates \mathcal{I} well* if the following holds for every pair of crossing – that is, not nested – separations $(A, B), (C, D) \in \mathcal{S}$:

For all $I_1, I_2 \in \mathcal{I}$ with $I_1 \subseteq A \cap C$ and $I_2 \subseteq B \cap D$ there is an $(E, F) \in \mathcal{S}$ such that $I_1 \subseteq E \subseteq A \cap C$ and $F \supseteq B \cup D$.

Note that such a separation satisfies both $(E, F) \leq (A, B)$ and $(E, F) \leq (C, D)$.

In our grid example, \mathcal{S} did not separate \mathcal{I} well, but we can mend this by adding to \mathcal{S} the four corner separations. And as soon as we do that, there is a nested subsystem that separates all four corners – for example, the set of the four corner separations.

More abstractly, the idea behind the notion of \mathcal{S} separating \mathcal{I} well is as follows. In the process of extracting \mathcal{N} from \mathcal{S} we may be faced with a pair of crossing separations (A, B) and (C, D) in \mathcal{S} that both separate two given sets $I_1, I_2 \in \mathcal{I}$, and wonder which of them to pick for \mathcal{N} . (Obviously we cannot choose both.) If \mathcal{S} separates \mathcal{I} well, however, we can avoid this dilemma by choosing (E, F) instead: this also separates I_1 from I_2 , and since it is nested with both (A, B) and (C, D) it will not prevent us from choosing either of these later too, if desired.

Let us call a separation $(E, F) \in \mathcal{S}$ *extremal* in \mathcal{S} if for all $(C, D) \in \mathcal{S}$ we have either $(E, F) \leq (C, D)$ or $(E, F) \leq (D, C)$. In particular, extremal separations are nested with all other separations in \mathcal{S} . Being extremal implies being \leq -minimal in \mathcal{S} ; if \mathcal{S} is nested, extremality and \leq -minimality are equivalent. If $(E, F) \in \mathcal{S}$ is extremal, then E is an \mathcal{S} -block; we call it an *extremal block* in \mathcal{S} .

A separation system, even a nested one, typically contains many extremal separations. For example, given a tree-decomposition of G with decomposition tree \mathcal{T} , the separations corresponding to the edges of \mathcal{T} that are incident with a leaf of \mathcal{T} are extremal in the (nested) set of all the separations of G corresponding to edges of \mathcal{T} .¹¹

Our next lemma shows that separating a set \mathcal{I} of \mathcal{S} -inseparable sets well is enough to guarantee the existence of an extremal separation among those that separate sets from \mathcal{I} . Call a separation \mathcal{I} -*relevant* if it weakly separates some two sets in \mathcal{I} . If all the separations in \mathcal{S} are \mathcal{I} -relevant, we call \mathcal{S} itself \mathcal{I} -*relevant*.

¹¹More precisely, every such edge of \mathcal{T} corresponds to an inverse pair of separations of which, usually, only one is extremal: the separation (A, B) for which A is the part V_t with t a leaf of \mathcal{T} . The separation (B, A) will not be extremal, unless $\mathcal{T} = K^2$.

Lemma 5.1. *Let \mathcal{R} be a separation system that is \mathcal{I} -relevant for some set \mathcal{I} of \mathcal{R} -inseparable sets. If \mathcal{R} separates \mathcal{I} well, then every \leq -minimal $(A, B) \in \mathcal{R}$ is extremal in \mathcal{R} . In particular, if $\mathcal{R} \neq \emptyset$ then \mathcal{R} contains an extremal separation.*

Proof. Consider a \leq -minimal separation $(A, B) \in \mathcal{R}$, and let $(C, D) \in \mathcal{R}$ be given. If (A, B) and (C, D) are nested, then the minimality of (A, B) implies that $(A, B) \leq (C, D)$ or $(A, B) \leq (D, C)$, as desired. So let us assume that (A, B) and (C, D) cross.

As (A, B) and (C, D) are \mathcal{I} -relevant and the sets in \mathcal{I} are \mathcal{R} -inseparable, we can find opposite corners of the cross-diagram $\{(A, B), (C, D)\}$ that each contains a set from \mathcal{I} . Renaming (C, D) as (D, C) if necessary, we may assume that these sets lie in $A \cap C$ and $B \cap D$, say $I_1 \subseteq A \cap C$ and $I_2 \subseteq B \cap D$. As \mathcal{R} separates \mathcal{I} well, there exists $(E, F) \in \mathcal{R}$ such that $I_1 \subseteq E \subseteq A \cap C$ and $F \supseteq B \cup D$, and hence $(E, F) \leq (A, B)$ as well as $(E, F) \leq (C, D)$. By the minimality of (A, B) , this yields $(A, B) = (E, F) \leq (C, D)$ as desired. \square

Let us say that a set \mathcal{S} of separations *distinguishes* two given \mathcal{S} -inseparable sets I_1, I_2 (or *distinguishes* them *properly*) if it contains a separation that separates them. If it contains a separation that separates them weakly, it *weakly distinguishes* I_1 from I_2 . We then also call I_1 and I_2 (*weakly*) *distinguishable* by \mathcal{S} , or (*weakly*) *\mathcal{S} -distinguishable*.

Here is our main result for this section:

Theorem 5.2. *Let \mathcal{S} be any separation system that separates some set \mathcal{I} of \mathcal{S} -inseparable sets of vertices well. Then \mathcal{S} has a nested \mathcal{I} -relevant subsystem $\mathcal{N}(\mathcal{S}, \mathcal{I}) \subseteq \mathcal{S}$ that weakly distinguishes all weakly \mathcal{S} -distinguishable sets in \mathcal{I} .*

Proof. If \mathcal{I} has no two weakly distinguishable elements, let $\mathcal{N}(\mathcal{S}, \mathcal{I})$ be empty. Otherwise let $\mathcal{R} \subseteq \mathcal{S}$ be the subsystem of all \mathcal{I} -relevant separations in \mathcal{S} . Then $\mathcal{R} \neq \emptyset$, and \mathcal{R} separates \mathcal{I} well. Let $\mathcal{E} \subseteq \mathcal{R}$ be the subset of those separations that are extremal in \mathcal{R} , and put

$$\bar{\mathcal{E}} := \{(A, B) \mid (A, B) \text{ or } (B, A) \text{ is in } \mathcal{E}\}.$$

By Lemma 5.1 we have $\bar{\mathcal{E}} \neq \emptyset$, and by definition of extremality all separations in $\bar{\mathcal{E}}$ are nested with all separations in \mathcal{R} . In particular, $\bar{\mathcal{E}}$ is nested.

Let

$$\mathcal{I}_{\mathcal{E}} := \{I \in \mathcal{I} \mid \exists (E, F) \in \mathcal{E} : I \subseteq E\}.$$

This is non-empty, since $\mathcal{E} \subseteq \mathcal{R}$ is non-empty and \mathcal{I} -relevant. Let us prove that \mathcal{E} weakly distinguishes all pairs of weakly distinguishable elements $I_1, I_2 \in \mathcal{I}$ with $I_1 \in \mathcal{I}_{\mathcal{E}}$. Pick $(A, B) \in \mathcal{R}$ with $I_1 \subseteq A$ and $I_2 \subseteq B$. Since $I_1 \in \mathcal{I}_{\mathcal{E}}$, there is an $(E, F) \in \mathcal{E}$ such that $I_1 \subseteq E$. By the extremality of (E, F) we have either $(E, F) \leq (A, B)$, in which case $I_1 \subseteq E$ and $I_2 \subseteq B \subseteq F$, or we have $(E, F) \leq (B, A)$, in which case $I_1 \subseteq E \cap A \subseteq E \cap F$. In both cases I_1 and I_2 are weakly separated by (E, F) .

As $\mathcal{I}' := \mathcal{I} \setminus \mathcal{I}_{\mathcal{E}}$ is a set of \mathcal{S} -inseparable sets with fewer elements than \mathcal{I} , induction gives us a nested \mathcal{I}' -relevant subsystem $\mathcal{N}(\mathcal{S}, \mathcal{I}')$ of \mathcal{S} that weakly

distinguishes all weakly distinguishable elements of \mathcal{I}' . Then

$$\mathcal{N}(\mathcal{S}, \mathcal{I}) := \overline{\mathcal{E}} \cup \mathcal{N}(\mathcal{S}, \mathcal{I}')$$

is \mathcal{I} -relevant and weakly distinguishes all weakly distinguishable elements of \mathcal{I} . As $\mathcal{I}' \subseteq \mathcal{I}$, and thus $\mathcal{N}(\mathcal{S}, \mathcal{I}') \subseteq \mathcal{R}$, the separations in $\overline{\mathcal{E}}$ are nested with those in $\mathcal{N}(\mathcal{S}, \mathcal{I}')$. Hence, $\mathcal{N}(\mathcal{S}, \mathcal{I})$ too is nested. \square

An important feature of the proof of Theorem 5.2 is that the subset $\mathcal{N}(\mathcal{S}, \mathcal{I})$ it constructs is *canonical*, given \mathcal{S} and \mathcal{I} : there are no choices made anywhere in the proof. We may thus think of \mathcal{N} as a recursively defined operator that assigns to every pair $(\mathcal{S}, \mathcal{I})$ as given in the theorem a certain nested subsystem $\mathcal{N}(\mathcal{S}, \mathcal{I})$ of \mathcal{S} . This subsystem $\mathcal{N}(\mathcal{S}, \mathcal{I})$ is canonical also in the structural sense that it is invariant under any automorphisms of G that leave \mathcal{S} and \mathcal{I} invariant.

To make this more precise, we need some notation. Every automorphism α of G acts also on (the set of) its vertex sets $U \subseteq V(G)$, on the collections \mathcal{X} of such vertex sets, on the separations (A, B) of G , and on the sets \mathcal{S} of such separations. We write U^α , \mathcal{X}^α , $(A, B)^\alpha$ and \mathcal{S}^α and so on for their images under α .

Corollary 5.3. *Let \mathcal{S} and \mathcal{I} be as in Theorem 5.2, and let $\mathcal{N}(\mathcal{S}, \mathcal{I})$ be the nested subsystem of \mathcal{S} constructed in the proof. Then for every automorphism α of G we have $\mathcal{N}(\mathcal{S}^\alpha, \mathcal{I}^\alpha) = \mathcal{N}(\mathcal{S}, \mathcal{I})^\alpha$. In particular, if \mathcal{S} and \mathcal{I} are invariant under the action of a group Γ of automorphisms of G , then so is $\mathcal{N}(\mathcal{S}, \mathcal{I})$.*

Proof. The proof of the first assertion is immediate from the construction of $\mathcal{N}(\mathcal{S}, \mathcal{I})$ from \mathcal{S} and \mathcal{I} . The second assertion follows, as

$$\mathcal{N}(\mathcal{S}, \mathcal{I})^\alpha = \mathcal{N}(\mathcal{S}^\alpha, \mathcal{I}^\alpha) = \mathcal{N}(\mathcal{S}, \mathcal{I})$$

for every $\alpha \in \Gamma$. \square

6 Separating the k -blocks of a graph

We now apply the theory developed in the previous sections to our original problem, of how to ‘decompose a graph G into its $(k+1)$ -connected components’. In the language of Section 5, we consider as \mathcal{S} the set of all proper k -separations of G , and as \mathcal{I} the set of its k -blocks. Our results from Section 5 rest on the assumption that the set \mathcal{R} of \mathcal{I} -relevant separations in \mathcal{S} separates \mathcal{I} well (Lemma 5.1). So the first thing we have to ask is: given crossing k -separations (A, B) and (C, D) such that $A \cap C$ and $B \cap D$ contain k -blocks b_1 and b_2 , respectively, is there a k -separation (E, F) such that $b_1 \subseteq E \subseteq A \cap C$?

If G is k -connected, there clearly is. Indeed, as the corners $A \cap C$ and $B \cap D$ each contain a k -block, they have order at least $k+1$, so their boundaries cannot have size less than k . But the sizes of these two corner boundaries sum to $|A \cap B| + |C \cap D| = 2k$, so they are both exactly k . We can thus take as (E, F) the corner separation $(A \cap C, B \cup D)$.

If G is not k -connected, we shall need another reason for these corner separations to have order at least k . This is a non-trivial problem. Our solution will be to assume inductively that those k -blocks that can be separated by a separation of order $\ell < k$ are already separated by such a separation selected earlier in the induction. Then the two corner separations considered above will have order at least k , since the k -blocks in the two corners are assumed not to have been separated earlier.

This approach differs only slightly from the more ambitious approach to build, inductively on ℓ , one nested set of separations which, for all ℓ at once, distinguishes every two ℓ -blocks by a separation of order at most ℓ . We shall construct an example showing that such a unified nested separation system need not exist. The subtle difference between our approach and this seemingly more natural generalization is that we use ℓ -separations for $\ell < k$ only with the aim to separate k -blocks; we do not aspire to separate all ℓ -blocks, including those that contain no k -block.

However we shall be able to prove that the above example is essentially the only one precluding the existence of a unified nested set of separations. Under a mild additional assumption saying that all blocks considered must be ‘robust’, we shall obtain one unified nested set of separations that distinguishes, for all ℓ simultaneously, all ℓ -blocks by a separation of order at most ℓ . All ℓ -blocks that have size at least $\frac{3}{2}\ell$ will be robust.

Once we have found our nested separation systems, we shall convert them into tree-decompositions as in Section 4. Both our separation systems and our tree-decompositions will be canonical in that they depend only on the structure of G . In particular, they will be invariant under the automorphism group $\text{Aut}(G)$ of G .

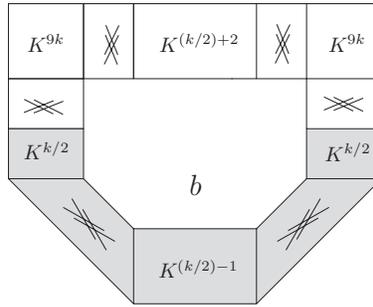


Figure 8: A horizontal k -separation needed to distinguish two k -blocks, crossed by a vertical $(k+1)$ -separation needed to distinguish two $(k+1)$ -blocks.

Let us now turn to our example showing that a graph need not have a ‘unified’ nested separation system \mathcal{N} of separations of mixed order that distinguishes, for every ℓ , distinct ℓ -blocks by a separation in \mathcal{N} of order at most ℓ . The graph depicted in Figure 8 arises from the disjoint union of a $K^{(k/2)-1}$, two $K^{k/2}$, a $K^{(k/2)+2}$ and two K^{9k} , by joining the $K^{(k/2)-1}$ completely to the two $K^{k/2}$, the $K^{(k/2)+2}$ completely to the two K^{9k} , the left $K^{k/2}$ completely to

the left K^{9k} , and the right $K^{k/2}$ completely to the right K^{9k} . The horizontal k -separator consisting of the two $K^{k/2}$ defines the only separation of order at most k that distinguishes the two k -blocks consisting of the top five complete graphs versus the bottom three. On the other hand, the vertical $(k+1)$ -separator consisting of the $K^{(k/2)-1}$ and the $K^{(k/2)+2}$ defines the only separation of order at most $(k+1)$ that distinguishes the two $(k+1)$ -blocks consisting, respectively, of the left $K^{k/2}$ and K^{9k} and the $K^{(k/2)+2}$, and of the right $K^{k/2}$ and K^{9k} and the $K^{(k/2)+2}$. Hence any separation system that distinguishes all k -blocks as well as all $(k+1)$ -blocks must contain both separations. Since the two separations cross, such a system cannot be nested.

In view of this example it may be surprising that we can find a separation system that distinguishes, for all $\ell \geq 0$ simultaneously, all *large* ℓ -blocks of G , those with at least $\lfloor \frac{3}{2}\ell \rfloor$ vertices. The example of Figure 8 shows that this value is best possible: here, all blocks are large except for the k -block b consisting of the two $K^{k/2}$ and the $K^{(k/2)-1}$, which has size $\frac{3}{2}k - 1$.

Indeed, we shall prove something considerably stronger: that the only obstruction to the existence of a unified tree-decomposition is a k -block that is not only not large but positioned exactly like b in Figure 8, inside the union of a k -separator and a larger separator crossing it.

Given integers k and K (where $k \leq K$ is the interesting case, but it is important formally to allow $k > K$), a k -inseparable set U is called *K -robust*¹² if for every k -separation (C, D) with $U \subseteq D$ and every separation (A, B) of order at most K such that $(A, B) \nparallel (C, D)$ and

$$|\partial(A \cap D)| < k > |\partial(B \cap D)|, \quad (15)$$

we have either $U \subseteq A$ or $U \subseteq B$. By $U \subseteq D$ and (15), the only way in which this can fail is that $|A \cap B| > k$ and U is contained in the union T of the boundaries of $A \cap D$ and $B \cap D$ (Fig. 9): exactly the situation of b in Figure 8.

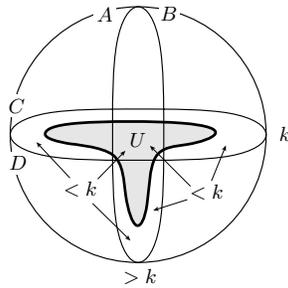


Figure 9: The shaded set U is k -inseparable but not K -robust.

It is obvious from the definition of robustness that

$$\text{for } k \geq K, \text{ every } k\text{-inseparable set is } K\text{-robust.} \quad (16)$$

¹²The parameter k is important here, too, but we suppress it for readability; it will always be stated explicitly in the context.

Let us call a k -inseparable set, in particular a k -block of G , *robust* if it is K -robust for every K (equivalently, for $K = |G|$). Our next lemma says that large k -blocks, those of size at least $\lfloor \frac{3}{2}k \rfloor$, are robust. But there are more kinds of robust sets than these: the vertex set of any K^{k+1} subgraph, for example, is a robust k -inseparable set.

Lemma 6.1. *Large k -blocks are robust.*

Proof. By the remark following the definition of ‘ K -robust’, it suffices to show that the set $T = \partial(A \cap D) \cup \partial(B \cap D)$ in Figure 9 has size at most $\frac{3}{2}k - 1$, regardless of the order of (A, B) . Let $\ell := |(A \cap B) \setminus C|$ be the size of the common link of the corners $A \cap D$ and $B \cap D$. By $|C \cap D| = k$ and (15) we have $2\ell \leq k - 2$, so $|T| = k + \ell \leq \frac{3}{2}k - 1$ as desired. \square

For the remainder of this paper, a *block* of G is again a subset of $V(G)$ that is a k -block for some k . The smallest k for which a block b is a k -block is its *rank*; let us denote this by $r(b)$. A block b that is given without a specified k is called *K -robust* if it is K -robust as an $r(b)$ -inseparable set. When we speak of a ‘robust k -block’ b , however, we mean the (stronger, see below) robustness as a k -inseparable set, not just as an $r(b)$ -inseparable set.

It is not difficult to find examples of K -robust blocks that are k -blocks but are not K -robust as a k -block, only as an ℓ -block for some $\ell < k$. A k -inseparable set that is K -robust as a k' -inseparable set for $k' > k$, however, is also K -robust as a k -inseparable set. More generally:

Lemma 6.2. *Let k, k' and K be integers.*

- (i) *Every k -inseparable set I containing a K -robust k' -inseparable set I' with $k \leq k'$ is K -robust.*
- (ii) *Every block b that contains a K -robust block b' is K -robust.*

Proof. (i) Suppose that I is not K -robust, and let this be witnessed by a k -separation (C, D) crossed by a separation (A, B) of order $m \leq K$. Put $S := C \cap D$ and $L := (A \cap B) \setminus C$. Then $I \subseteq S \cup L$, as remarked after the definition of ‘ K -robust’.

Extend S into L to a k' -set S' that is properly contained in $S \cup L$ (which is large enough, since it contains $I' \subseteq I$), and put $C' := C \cup S'$. Then (C', D) is a k' -separation with separator S' and corners $D \cap A$ and $D \cap B$ with (A, B) , whose boundaries by assumption have size less than $k \leq k'$. As I' is K -robust, it lies in one of these corners, say $I' \subseteq A \cap D$. Since

$$|I'| > k' \geq k > |\partial(A \cap D)|,$$

this implies that I' has a vertex in the interior of the corner $A \cap D$. As $I' \subseteq I$, this contradicts the fact that $I \subseteq S \cup L$.

(ii) The block b is an $r(b)$ -inseparable set containing the K -robust $r(b')$ -inseparable set b' . If $b = b'$ then $r(b) = r(b')$. If $b \supsetneq b'$, then b' is not maximal as an ℓ -inseparable set for any $\ell \leq r(b)$, giving $r(b') > r(b)$. Hence $r(b) \leq r(b')$ either way, so b is a K -robust block by (i). \square

Let us call two blocks *distinguishable* if neither contains the other. It is not hard to show that distinguishable blocks b_1, b_2 can be separated in G by a separation of order $r \leq \min\{r(b_1), r(b_2)\}$. We denote the smallest such r by

$$\kappa(b_1, b_2) \leq \min\{r(b_1), r(b_2)\},$$

and say that b_1 and b_2 are k -*distinguishable* for a given integer k if $\kappa(b_1, b_2) \leq k$. Note that distinct k -blocks are k -distinguishable, but they might also be ℓ -distinguishable for some $\ell < k$.

A set \mathcal{S} of separations *distinguishes* two k -blocks if it contains a separation of order at most k that separates them. It *distinguishes* two blocks b_1, b_2 given without a specified k if it contains a separation of order $r \leq \min\{r(b_1), r(b_2)\}$ that separates them.¹³ If \mathcal{S} contains a separation of order $\kappa(b_1, b_2)$ that separates two blocks or k -blocks b_1, b_2 , we say that \mathcal{S} distinguishes them *efficiently*.

Theorem 6.3. *For every finite graph G and every integer $k \geq 0$ there is a tight, nested, and $\text{Aut}(G)$ -invariant separation system \mathcal{N}_k that distinguishes every two k -distinguishable k -robust blocks efficiently. In particular, \mathcal{N}_k distinguishes every two k -blocks efficiently.*

Proof. Let us rename the integer k given in the theorem as K . Recursively for all integers $0 \leq k \leq K$ we shall construct a sequence of separation systems \mathcal{N}_k with the following properties:

- (i) \mathcal{N}_k is tight, nested, and $\text{Aut}(G)$ -invariant;
- (ii) $\mathcal{N}_{k-1} \subseteq \mathcal{N}_k$ (put $\mathcal{N}_{-1} := \emptyset$);
- (iii) every separation in $\mathcal{N}_k \setminus \mathcal{N}_{k-1}$ has order k ;
- (iv) \mathcal{N}_k distinguishes every two K -robust k -blocks.
- (v) every separation in $\mathcal{N}_k \setminus \mathcal{N}_{k-1}$ separates some K -robust k -blocks that are not distinguished by \mathcal{N}_{k-1} .

We claim that \mathcal{N}_K will satisfy the assertions of the theorem for $k = K$. Indeed, consider two K -distinguishable K -robust blocks b_1, b_2 . Then

$$\kappa := \kappa(b_1, b_2) \leq \min\{K, r(b_1), r(b_2)\},$$

so b_1, b_2 are κ -inseparable and extend to distinct κ -blocks b'_1, b'_2 . These are again K -robust, by Lemma 6.2 (i). Hence by (iv), $\mathcal{N}_\kappa \subseteq \mathcal{N}_K$ distinguishes $b'_1 \supseteq b_1$ from $b'_2 \supseteq b_2$, and it does so efficiently by definition of κ .

It remains to construct the separation systems \mathcal{N}_k .

Let $k \geq 0$ be given, and assume inductively that we already have separation systems $\mathcal{N}_{k'}$ satisfying (i)–(v) for $k' = 0, \dots, k-1$. (For $k = 0$ we have nothing

¹³Unlike in the definition just before Theorem 5.2, we no longer require that the blocks we wish to separate be \mathcal{S} -inseparable for the entire set \mathcal{S} .

but the definition of $\mathcal{N}_{-1} := \emptyset$, which has $V(G)$ as its unique \mathcal{N}_{-1} -block.) Let us show the following:

For all $0 \leq \ell \leq k$, any two K -robust ℓ -blocks b_1, b_2 that are not distinguished by $\mathcal{N}_{\ell-1}$ satisfy $\kappa(b_1, b_2) = \ell$. (17)

This is trivial for $\ell = 0$; let $\ell > 0$. If $\kappa(b_1, b_2) < \ell$, then the $(\ell - 1)$ -blocks $b'_1 \supseteq b_1$ and $b'_2 \supseteq b_2$ are distinct. By Lemma 6.2 (i) they are again K -robust. Thus by hypothesis (iv) they are distinguished by $\mathcal{N}_{\ell-1}$, and hence so are b_1 and b_2 , contrary to assumption.

By hypothesis (iii), every k -block is \mathcal{N}_{k-1} -inseparable, so it extends to some \mathcal{N}_{k-1} -block; let \mathcal{B} denote the set of those \mathcal{N}_{k-1} -blocks that contain more than one K -robust k -block. For each $b \in \mathcal{B}$ let \mathcal{I}_b be the set of all K -robust k -blocks contained in b . Let \mathcal{S}_b denote the set of all those k -separations of G that separate some two elements of \mathcal{I}_b and are nested with all the separations in \mathcal{N}_{k-1} .

Clearly \mathcal{S}_b is symmetric and the separations in \mathcal{S}_b are proper (since they distinguish two k -blocks), so \mathcal{S}_b is a separation system of G . By (17) for $\ell = k$, the separations in \mathcal{S}_b are tight. Our aim is to apply Theorem 5.2 to extract from \mathcal{S}_b a nested subsystem \mathcal{N}_b that we can add to \mathcal{N}_{k-1} .

Before we verify the premise of Theorem 5.2, let us prove that it will be useful: that the nested separation system $\mathcal{N}_b \subseteq \mathcal{S}_b$ it yields can distinguish¹⁴ all the elements of \mathcal{I}_b . This will be the case only if \mathcal{S}_b does so, so let us prove this first:

(*) \mathcal{S}_b distinguishes every two elements of \mathcal{I}_b .

For a proof of (*) we have to find for any two k -blocks $I_1, I_2 \in \mathcal{I}_b$ a separation in \mathcal{S}_b that separates them. Applying Lemma 2.1 with the set \mathcal{S} of all separations of order at most k , we can find a separation $(A, B) \in \mathcal{S}$ such that $I_1 \subseteq A$ and $I_2 \subseteq B$. Choose (A, B) so that it is nested with as many separations in \mathcal{N}_{k-1} as possible. We prove that $(A, B) \in \mathcal{S}_b$, by showing that (A, B) has order exactly k and is nested with every separation $(C, D) \in \mathcal{N}_{k-1}$. Let $(C, D) \in \mathcal{N}_{k-1}$ be given.

Being elements of \mathcal{I}_b , the sets I_1 and I_2 cannot be separated by fewer than k vertices, by (17). Hence (A, B) has order exactly k . Since I_1 is k -inseparable it lies on one side of (C, D) , say in C , so $I_1 \subseteq A \cap C$. As (C, D) does not separate I_1 from I_2 , we then have $I_2 \subseteq B \cap C$.

Let $\ell < k$ be such that $(C, D) \in \mathcal{N}_\ell \setminus \mathcal{N}_{\ell-1}$. By hypothesis (v) for ℓ , there are K -robust ℓ -blocks $J_1 \subseteq C$ and $J_2 \subseteq D$ that are not distinguished by $\mathcal{N}_{\ell-1}$. By (17),

$$\kappa(J_1, J_2) = \ell. \quad (18)$$

Let us show that we may assume the following:

The corner separations of the corners $A \cap C$ and $B \cap C$ are nested with every separation $(C', D') \in \mathcal{N}_{k-1}$ that (A, B) is nested with. (19)

¹⁴As the elements of \mathcal{I}_b are k -blocks, we have two notions of ‘distinguish’ that could apply: the definition given before Theorem 5.2, or that given before Theorem 6.3. However, as \mathcal{S}_b consists of k -separations and all the elements of \mathcal{I}_b are \mathcal{S}_b -inseparable, the two notions coincide.

Since (C, D) and (C', D') are both elements of \mathcal{N}_{k-1} , they are nested with each other. Thus,

$$(A, B) \parallel (C', D') \parallel (C, D).$$

Unless (A, B) is nested with (C, D) (in which case our proof of $(*)$ is complete), this implies by Lemma 2.2 that (C', D') is nested with all the corner separations of the cross-diagram for (A, B) and (C, D) , especially with those of the corners $A \cap C$ and $B \cap C$ that contain I_1 and I_2 . This proves (19).

Since the corner separations of $A \cap C$ and $B \cap C$ are nested with the separation $(C, D) \in \mathcal{N}_{k-1}$ that (A, B) is not nested with (as we assume), (19) and the choice of (A, B) imply that

$$|\partial(A \cap C)| \geq k + 1 \quad \text{and} \quad |\partial(B \cap C)| \geq k + 1.$$

Since the sizes of the boundaries of two opposite corners sum to

$$|A \cap B| + |C \cap D| = k + \ell,$$

this means that the boundaries of the corners $A \cap D$ and $B \cap D$ have sizes $< \ell$. Since J_2 is K -robust as an ℓ -block, we thus have $J_2 \subseteq A \cap D$ or $J_2 \subseteq B \cap D$, say the former. But as $J_1 \subseteq C \subseteq B \cup C$, this contradicts (18), completing the proof of $(*)$.

Let us now verify the premise of Theorem 5.2:

$(**)$ \mathcal{S}_b separates \mathcal{I}_b well.

Consider a pair $(A, B), (C, D) \in \mathcal{S}_b$ of crossing separations with sets $I_1, I_2 \in \mathcal{I}_b$ such that $I_1 \subseteq A \cap C$ and $I_2 \subseteq B \cap D$. We shall prove that $(A \cap C, B \cup D) \in \mathcal{S}_b$.

By (17) and $I_1, I_2 \in \mathcal{I}_b$, the boundaries of the corners $A \cap C$ and $B \cap D$ have size at least k . Since their sizes sum to $|A \cap B| + |C \cap D| = 2k$, they each have size exactly k . Hence $(A \cap C, B \cup D)$ has order k and is nested with every separation $(C', D') \in \mathcal{N}_{k-1}$ by Lemma 2.2, because $(A, B), (C, D) \in \mathcal{S}_b$ implies that (A, B) and (C, D) are both nested with $(C', D') \in \mathcal{N}_{k-1}$. This completes the proof of $(**)$.

By $(*)$ and $(**)$, Theorem 5.2 implies that \mathcal{S}_b has a nested \mathcal{I}_b -relevant subsystem $\mathcal{N}_b := \mathcal{N}(\mathcal{S}_b, \mathcal{I}_b)$ that weakly distinguishes all the sets in \mathcal{I}_b . But these are k -inseparable and hence of size $> k$, so they cannot lie inside a k -separator. So \mathcal{N}_b even distinguishes the sets in \mathcal{I}_b properly. Let

$$\mathcal{N}_{\mathcal{B}} := \bigcup_{b \in \mathcal{B}} \mathcal{N}_b \quad \text{and} \quad \mathcal{N}_k := \mathcal{N}_{k-1} \cup \mathcal{N}_{\mathcal{B}}.$$

Let us verify the inductive statements (i)–(v) for k . We noted earlier that every \mathcal{S}_b is tight, hence so is every \mathcal{N}_b . The separations in each \mathcal{N}_b are nested with each other and with \mathcal{N}_{k-1} . Separations from different sets \mathcal{N}_b are nested

by Lemma 2.3. So the entire set \mathcal{N}_k is nested. Since \mathcal{N}_{k-1} is $\text{Aut}(G)$ -invariant, by hypothesis (i), so is \mathcal{B} . For every automorphism α and every $b \in \mathcal{B}$ we then have $\mathcal{I}_{b^\alpha} = (\mathcal{I}_b)^\alpha$ and $\mathcal{S}_{b^\alpha} = (\mathcal{S}_b)^\alpha$, so Corollary 5.3 yields $(\mathcal{N}_b)^\alpha = \mathcal{N}_{b^\alpha}$. Thus, $\mathcal{N}_{\mathcal{B}}$ is $\text{Aut}(G)$ -invariant too, completing the proof of (i). Assertions (ii) and (iii) hold by definition of \mathcal{N}_k . Assertion (iv) is easy too: if two K -robust k -blocks are not distinguished by \mathcal{N}_{k-1} they will lie in the same \mathcal{N}_{k-1} -block b , and hence be distinguished by \mathcal{N}_b . Assertion (v) holds, because each \mathcal{N}_b is \mathcal{I}_b -relevant. \square

Let us call two blocks b_1, b_2 of G *robust* if there exists a k for which they are robust k -blocks.¹⁵ For $k = |G|$, Theorem 6.3 then yields our ‘unified’ nested separation system that separates all robust blocks by a separation of the lowest possible order:

Corollary 6.4. *For every finite graph G there is a tight, nested, and $\text{Aut}(G)$ -invariant separation system \mathcal{N} that distinguishes every two distinguishable robust blocks efficiently.* \square

Let us now turn the separation systems \mathcal{N}_k of Theorem 6.3 and its proof into tree-decompositions:

Theorem 6.5. *For every finite graph G and every integer K there is a sequence $(\mathcal{T}_k, \mathcal{V}_k)_{k \leq K}$ of tree-decompositions such that, for all $k \leq K$,*

- (i) *every k -inseparable set is contained in a unique part of $(\mathcal{T}_k, \mathcal{V}_k)$;*
- (ii) *distinct K -robust k -blocks lie in different parts of $(\mathcal{T}_k, \mathcal{V}_k)$;*
- (iii) *$(\mathcal{T}_k, \mathcal{V}_k)$ has adhesion at most k ;*
- (iv) *if $k > 0$ then $(\mathcal{T}_{k-1}, \mathcal{V}_{k-1}) \preceq (\mathcal{T}_k, \mathcal{V}_k)$;*
- (v) *$\text{Aut}(G)$ acts on \mathcal{T}_k as a group of automorphisms.*

Proof. Consider the nested separation system \mathcal{N}_K given by Theorem 6.3. As in the proof of that theorem, let \mathcal{N}_k be the subsystem of \mathcal{N}_K consisting of its separations of order at most k . By Theorem 6.3, \mathcal{N}_K is $\text{Aut}(G)$ -invariant, so this is also true for all \mathcal{N}_k with $k < K$.

Let $(\mathcal{T}_k, \mathcal{V}_k)$ be the tree-decomposition associated with \mathcal{N}_k as in Section 4. Then (v) holds by Corollary 3.5, (iii) and (iv) by Theorem 4.8 (iii) and (iv). By (iii) and [3, Lemma 12.3.1], any k -inseparable set is contained in a unique part of $(\mathcal{T}_k, \mathcal{V}_k)$, giving (i). By (iv) in the proof of Theorem 6.3, \mathcal{N}_k distinguishes every two K -robust k -blocks, which implies (ii) by (i) and Theorem 4.8 (iii). \square

From Theorem 6.5 we can finally deduce the two results announced in the Introduction, Theorems 1 and 2.

Theorem 1 follows by taking as K the integer k given in Theorem 1, and then considering the decomposition $(\mathcal{T}_k, \mathcal{V}_k)$ for $k = K$. Indeed, consider two

¹⁵By Lemma 6.2 (i), this is equivalent to saying that they are robust $r(b_i)$ -blocks, that is, K -robust $r(b_i)$ -blocks for $K = |G|$.

k -blocks b_1, b_2 that Theorem 1 claims are distinguished efficiently by $(\mathcal{T}_k, \mathcal{V}_k)$. By Theorem 6.5 (ii), b_1 and b_2 lie in different parts of $(\mathcal{T}_k, \mathcal{V}_k)$. Let $k' := \kappa(b_1, b_2) \leq k$. By Lemma 6.2 (i), the k' -blocks $b'_1 \supseteq b_1$ and $b'_2 \supseteq b_2$ are again K -robust. Hence by Theorem 6.5 (ii) for k' , they lie in different parts of $(\mathcal{T}_{k'}, \mathcal{V}_{k'})$. Consider an adhesion set of $(\mathcal{T}_{k'}, \mathcal{V}_{k'})$ on the path in $\mathcal{T}_{k'}$ between these parts. By Theorem 6.5 (iii), this set has size at most k' , and by Theorem 6.5 (iv) it is also an adhesion set of $(\mathcal{T}_k, \mathcal{V}_k)$ between the two parts of $(\mathcal{T}_k, \mathcal{V}_k)$ that contain b_1 and b_2 .

Theorem 2 follows from Theorem 6.5 for $K = |G|$; recall that robust k -blocks are K -robust for $K = |G|$.

7 Outlook

There are two types of question that arise from the context of this paper, but which we have not addressed.

The first of these concerns its algorithmic aspects. How hard is it

- to decide whether a given graph has a k -block;
- to find all the k -blocks in a given graph;
- to compute the canonical tree-decompositions whose existence we have shown?

Note that our definitions leave some leeway in answering the last question. For example, consider a graph G that consists of two disjoint complete graphs K, K' of order 10 joined by a long path P . For $k = 5$, this graph has only two k -blocks, K and K' . One tree-decomposition of G that is invariant under its automorphisms has as parts the graphs K, K' and all the K_2 s along the path P , its decomposition tree again being a long path. This tree-decomposition is particularly nice also in that it also distinguishes the ℓ -blocks of G not only for $\ell = k$ but for all ℓ such that G has an ℓ -block, in particular, for $\ell = 1$.

However if we are only interested in k -blocks for $k = 5$, this decomposition can be seen as unnecessarily fine in that it has many parts containing no k -block. We might, in this case, prefer a tree-decomposition that has only two parts, and clearly there is such a tree-decomposition that is invariant under $\text{Aut}(G)$, of adhesion 1 or 2 depending on the parity of $|P|$.

This tree-decomposition, however, is suboptimal in yet another respect: we might prefer decompositions in which any part that does contain a k -block contains nothing but this k -block. Our first decomposition satisfies this, but there is another that does too while having fewer parts: the path-decomposition into three parts whose middle part is P and whose leaf parts are K and K' .

We shall look at these possibilities and associated algorithms in more detail in [1]. However we shall not make an effort to optimize these algorithms from a complexity point of view, so the above three questions will be left open.

Since our tree-decompositions are canonical, another obvious question is whether they, or refinements, can be used to tackle the graph isomorphism problem. Are there natural classes of graphs for which we can

- describe the parts of our canonical tree-decompositions in more detail;
- use this to decide graph isomorphism for such classes in polynomial time?

Another broad question that we have not touched upon, not algorithmic, is the following. Denote by $\beta(G)$ the greatest integer k such that G has a k -block (or equivalently: has a k -inseparable set of vertices). This seems to be an interesting graph invariant; for example, in a network G one might think of the nodes of a $\beta(G)$ -block as locations to place some particularly important servers that should still be able to communicate with each other when much of the network has failed.

From a mathematical point of view, it seems interesting to ask how β interacts with other graph invariants. For example, what average degree will force a graph to contain a k -block for given k ? What can we say about the structure of graphs that contain no k -block but have large tree-width?

Some preliminary results in this direction are obtained in [2], but even for the questions we address we do not have optimal results.

Acknowledgement

Many ideas for this paper have grown out of an in-depth study of the treatise [6] by Dunwoody and Krön, which we have found both enjoyable and inspiring.

References

- [1] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs. in preparation.
- [2] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Sufficient conditions for the existence of a k -block in a finite graph. in preparation.
- [3] R. Diestel. *Graph Theory*. Springer, 4th edition, 2010.
- [4] R. Diestel, K.Yu. Gorbunov, T. Jensen, and C. Thomassen. Highly connected sets and the excluded grid theorem. *J. Combin. Theory (Series B)*, 75:61–73, 1999.
- [5] M.J. Dunwoody. Cutting up graphs. *Combinatorica*, 2:15–23, 1982.
- [6] M.J. Dunwoody and B. Krön. Vertex cuts. arXiv:0905.0064, 2009.
- [7] F. Hundertmark. Profiles. An algebraic approach to combinatorial connectivity. arXiv:1110.6207, 2011.
- [8] W. Mader. Über n -fach zusammenhängende Eckenmengen in Graphen. *J. Combin. Theory (Series B)*, 25:74–93, 1978.
- [9] B.A. Reed. Tree width and tangles: a new connectivity measure and some applications. In R.A. Bailey, editor, *Surveys in Combinatorics*. Cambridge Univ. Press, 1997.
- [10] N. Robertson and P.D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)*, 52:153–190, 1991.
- [11] W. T. Tutte. *Graph Theory*. Addison-Wesley, 1984.