The tree-like connectivity structure of finite graphs and matroids

Dissertation
zur Erlangung des Doktorgrades
der Fakultät für Mathematik, Informatik
und Naturwissenschaften
der Universität Hamburg

vorgelegt

im
Fachbereich Mathematik

von

Fabian Hundertmark

aus Hamburg

Hamburg
2013

Als Dissertation angenommen vom Fachbereich
Mathematik der Universität Hamburg

aufgrund der Gutachten von Prof. Dr. Reinhard Diestel
und Prof. Dr. Mathias Schacht

Hamburg, den 25.09.2012

Prof. Dr. Ulf Kühn
Leiter des Fachbereichs Mathematik

Acknowledgement

I am very grateful to my supervisor Reinhard Diestel for all his trust and support. I would like to thank all my colleagues in the workgroup for various discussions and refreshing coffee breaks. I thank Matthias Hamann not only for taking such good care for our wonderful office plant but also for his various encouragements and productive discussions. I also thank Maya Stein, for sharing a lot of mandarins during her stay in Hamburg, and for inviting me to Santiago de Chile afterwards. Finally I want to thank my close friends and my family, och min älskad Janna.

Contents

Chapter 1

Introduction

All graph-theoretic terms not defined within this thesis are used as in [5]. A *graph* is a pair $G = (V, E)$, where V is an arbitrary set and $E \subseteq [V]^2$ is a set of two-element subsets of V. Without loss of generality we may assume V and E to be disjoint. An element $v \in V$ is called a *vertex of* G whereas an element $e \in E$ is called an *edge of* G. We refer to the set of vertices of a graph G by $V(G)$ and to the set of edges by $E(G)$. A graph H is a *subgraph* of a graph G, denoted by $H \subseteq G$, if we have $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$; if we have $E(H) = E(G) \cap [V(H)]^2$ then H is called an *induced subgraph* of G. For $A \subseteq V(G)$ we let $G[A]$ denote the unique induced subgraph $H \subseteq G$ with $V(H) = A$, and we write $G - A$ for $G[V(G) \setminus A]$ and we abbreviate $G - \{v\}$ to $G - v$. Furthermore, for a subset $F \subseteq E$ we write $G - F$ for the graph $(V, E \setminus F)$ and $G - e$ instead of $G - \{e\}$. A graph is called *trivial* if it consists of at most one vertex. The (unique) graph which has no vertices and no edges is called the *empty graph*. The *union* $H_1 \cup H_2$ of two graphs H_1, H_2 is defined as $H_1 \cup H_2 := \big(V(H_1) \cup V(H_2), E(H_1) \cup E(H_2)\big)$; the *intersection* $H_1 \cap H_2$ is defined analogously as $H_1 \cap H_2 := \big(V(H_1) \cap V(H_2), E(H_1) \cap E(H_2)\big)$. Two graphs are called *disjoint* if their intersection is the empty graph.

A *path (of length k)* in G is a subgraph $P \subseteq G$ of the form $V(P) = \{v_0, \ldots, v_k\}$ where $k \geq 0$ is an integer and for $i \neq j \in \{0, \ldots, k\}$ we have $v_i \neq v_j$ and $\{v_i, v_j\} \in E(P) \Leftrightarrow |j - i| = 1$. The vertices v_0 and v_k are called the *end-vertices* of P, which P is said to *join*; all other vertices are called its *inner* vertices. A graph G is called *connected* if every two of its vertices are joined by a path in G. Given an integer k a graph is called *k-connected* if it has at least $k + 1$ vertices and $G - S$ is connected for every set $S \subseteq V$ with $|S| < k$. The maximal integer k for which G is k-connected is called the *connectivity of* G, denoted by $\kappa(G)$. Notice that a non-trivial graph is connected if and only if it is 1-connected.

Consider a graph G of connectivity 0. What do we know about its 'connectivity structure', other than being not connected? We do know that G is the disjoint union of its maximal connected subgraphs, or its *connected components*. This is quite easy to see if one realizes that 'being joined by a path in G' forms an equivalence relation on the set of vertices of G[1]. That is, for further investigations of the connectivity structure of G, we may investigate each of its

[1]Each vertex v is joined to itself by the trivial path $(\{v\}, \emptyset)$.

connected components independently, and may therefore as well assume that G itself is connected.

Now consider a graph G with $\kappa(G) = 1$. A natural question which then arises is whether or not G contains some 2-connected subgraph. It is easy to see that this is equivalent to asking whether or not G contains a *cycle*, which is a graph obtained by a path of length at least 2 by adding an edge between its two endvertices. A graph which contains no cycle is called a *forest*, and a connected forrest is called a *tree*. Now let us assume that G is not a tree and consider its maximal 2-connected subgraphs. These are not necessarily mutually disjoint. (We may for example consider a graph which is the union of two cycles which intersect in a unique common vertex.) But they may pairwise intersect in at most one vertex: if two 2-connected subgraphs have more than one vertex in common, then their union is obviously also 2-connected. The maximal 2-connected subgraphs are in fact connected in a 'tree-like way'. To make this precise we first have to introduce the notion of a block.

A *bridge* is an edge $e = \{v, w\}$ such that v and w lie in distinct components of $G-e$. Now a *block of* G is a set of vertices which either is the set of vertices of a maximal 2-connected subgraph, a bridge, or a singleton consisting of a vertex which has no neighbors in G. A vertex v is called a *cutvertex* (or 1-*separator*) if $G-v$ has more components then G does. There is the following relation between the blocks and cutvertices of a graph: each cutvertex is the intersection of two or more blocks, and the intersection of any set of blocks is either a cutvertex or empty. Let $\mathcal{B}(G)$ denote the set of blocks and let $V_C(G)$ be the set of cutvertices of G. Then we define the *block-graph of* G as the graph with $\mathcal{B}(G) \cup V_C(G)$ as its set of vertices, an edge between each block b and cutvertex v such that $v \in b$, and no other edges. The following statement is a well-known fact, for a proof see for example [5, Lemma 3.1.4].

The block-graph of a connected graph is a tree. (1.1)

This clearly clarifies (and strengthens) the rather vague statement from above, that the maximal 2-connected subgraphs of a graph are connected in a tree-like way. We call the block-graph of a connected graph its *block-cutvertex tree*. A graph together with its block-cutvertex tree is depicted in Figure 1.1.

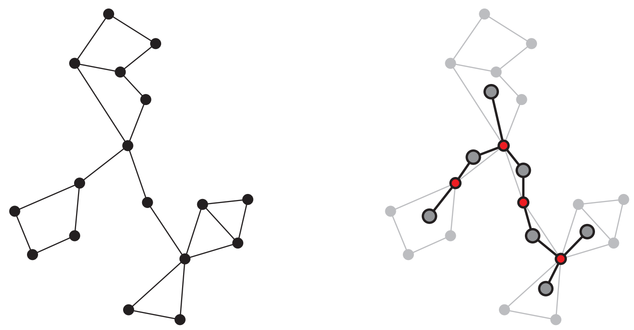

(a) *(b)*

Figure 1.1: A graph *(a)* and its block-cutvertex tree*(b)*.

While the information "$\kappa(G) = 1$" just tells us that G is connected and that there is at least one cutvertex in G, the block-cutvertex tree encapsulates much more: it describes the global layout of G with respect to its cutvertices.

The standard way in modern graph theory to describe such a tree-like global layout (with respect to small-order separators) is the notion of a tree-decomposition (of small adhesion), as defined by Robertson and Seymour in their well-known series on Graph Minor Theory (for an introduction see Diestel [5, Chapter 12]). As in [5] we define a *tree-decomposition* of G as a pair $(\mathcal{T}, \mathcal{V})$ of a tree $\mathcal{T}$ and a family $\mathcal{V} = (V_t)_{t \in \mathcal{T}}$ of subsets $V_t \subseteq V(G)$, one for every node[2] of $\mathcal{T}$, such that[3]:

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

(T2) $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever t_2 lies on the t_1–t_3 path in $\mathcal{T}$;

(T3) for every edge $e \in G$ there exists a $t \in \mathcal{T}$ such that both ends of e lie in V_t.

Two *parts* V_s, $V_t \in \mathcal{V}$ are *adjacent* if s and t are adjacent in $\mathcal{T}$. The intersections of adjacent parts are the *adhesion-sets* of $(\mathcal{T}, \mathcal{V})$ and the *adhesion* of $(\mathcal{T}, \mathcal{V})$ is the maximum cardinality of its adhesion-sets. A *torso* of $(\mathcal{T}, \mathcal{V})$ is a graph H obtained from a subgraph $G[V_t]$ of G which is induced on a part V_t of $(\mathcal{T}, \mathcal{V})$, where in addition in H each pair of vertices which is contained in a common adhesion set is joined by an edge (if they were already we simply keep this edge).

With this notion of a tree-decomposition at hand we can easily describe Tutte's approach from the 1960's to describe the connectivity structure of a 2-connected graph (cp. [17]):

Every graph of connectivity 2 *has a tree-decomposition of adhesion* 2 *all whose torsos each are either* 3*-connected or a cycle.* (1.2)

To add some edges to a graph will, in general, change the connectivity structure of that graph. However, for the additional torso edges in the tree-decomposition of (1.2) this is not the case: if in the construction of the torso obtained from a part V_t we add an edge between two vertices that lie in the same adhesion-set $V_t \cap V_s$, say, then those vertices are joined in G by a path that is contained in the union over all parts $V_{t'}$ such that t' lies in the component of $\mathcal{T} - t$ in which s lies. Hence, a set of independent paths in a torso (of this particular tree-decomposition) gives rise to a set of independent paths in the whole graph G with the same endvertices.

At this point we need a classical result in graph theory which sometimes is referred to as the 'global version' of Menger's Theorem (cp. Diestel [5, Theorem 3.3.6]):

A graph is k-connected if and only if it contains k independent paths between any two vertices. (1.3)

So those parts of the tree-decomposition given by (1.2) which give rise to 3-connected torsos, induce (maximal) subgraphs any two vertices of which are joined in G by at least 3 independent paths. However, such a subgraph, say H,

[2] As in [4], we refer to the vertices of our decomposition trees as its *nodes*.

[3] In order to stay consistent with Section 2 we changed the order of axioms (T2) and (T3) of the definition in [5].

does not need to be 3-connected, since those independent paths found in G may require vertices of $G - H$. But still, those subgraphs H should somehow be considered as the '3-connected pieces' of G.

Then (1.2) can be interpreted as follows: every 2-connected graph can be decomposed in a tree-like way into its 3-connected pieces. This gave rise to the following question:

Given an integer $k \geq 3$, can we decompose a k-connected graph in a tree-like way into its $(k+1)$-connected pieces? (1.4)

But a suitable general notion of '$(k+1)$-connected piece' was still to be found. Unfortunately, for $k \geq 3$ the concept of a torso is not suitable to model the notion of a $(k+1)$-connected piece. Once the adhesion-sets are allowed to consist of more than two vertices, a set of independent paths in a torso need not give rise to a set of independent paths in G (even if the adhesions-sets form minimal separators). However, we can avoid this problem by using the property of 'being joined by at least 3 independent paths in G' directly. In fact, we may use a slightly weaker property: we consider sets of vertices which are not separable in G by fewer than 3 vertices.

More general: given an integer k, we call a set I of at least $k+1$ vertices *k-inseparable* if it cannot be separated in G by at most k vertices. Now a *k-block* is defined as a maximal k-inseparable set of vertices. Then it is easy to see that the 0-blocks of a graph G are precisely its connected components while the 1-blocks of G are precisely its blocks (as defined above). And for $\kappa(G) = 2$, the 2-blocks of G are identified by the tree-decomposition given by (1.2): they are precisely those parts which give rise to a torso that is either 3-connected or a complete graph on three vertices (which is a cycle of length 3).

It is obvious from the definition of a k-block that distinct k-blocks, b_1 and b_2 say, can be separated by a set S of at most k vertices; let us denote the minimum size of such a set S by $\kappa(b_1, b_2)$. We say that a tree-decomposition $(\mathcal{T}, \mathcal{V})$ *distinguishes* the k-blocks b_1 and b_2 if there are nodes $t_1 \neq t_2 \in \mathcal{T}$ such that $b_i \subseteq V_{t_i}$ and there are nodes s, t on the unique path in $\mathcal{T}$ between t_1 and t_2 such that $|V_s \cap V_t| \leq k$; if we have $|V_s \cap V_t| = \kappa(b_1, b_2)$ then $(\mathcal{T}, \mathcal{V})$ *efficiently* distinguishes b_1 and b_2.

In Section 3 of this thesis we present our solution given in [4] to the problem stated in (1.4), using the notion of a k-block to identify the $(k+1)$-connected pieces of a graph. An essential property of a tree-decomposition $(\mathcal{T}, \mathcal{V})$ which we construct for a graph G is that it is invariant under the automorphisms of G: every automorphism α of G induces an automorphism α' of $\mathcal{T}$ such that for every $t \in \mathcal{T}$ and every $v \in V(G)$ we have $v \in V_t \Leftrightarrow \alpha(v) \in V_{\alpha'(t)}$. We show:

Theorem 1.1. *Given any integer $k \geq 0$, every finite graph G has an* $\mathrm{Aut}(G)$*-invariant tree-decomposition of adhesion at most k that efficiently distinguishes all its k-blocks.*

As we have pointed out in [4], we can combine those tree-decompositions found for different values of k into one overall tree-decomposition if, for each k, we consider only 'robust' k-blocks (a rather technical condition on a k-block which, however, is satisfied by most k-blocks, in particular by those that are large enough in terms of k). We are able to do so by constructing the corresponding tree-decomposition for a $k > 0$ as a refinement of the tree-decomposition already constructed for $k-1$. Here we say that a tree-decomposition $(\mathcal{T}_m, \mathcal{V}_m)$ *is refined by* a tree-decomposition $(\mathcal{T}_n, \mathcal{V}_n)$, denoted by $(\mathcal{T}_m, \mathcal{V}_m) \preccurlyeq (\mathcal{T}_n, \mathcal{V}_n)$, 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 $\mathcal{T}_m$. We then obtain the following result:

Theorem 1.2. *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) $(\mathcal{T}_k, \mathcal{V}_k) \preccurlyeq (\mathcal{T}_{k+1}, \mathcal{V}_{k+1})$*;*

(iii) $(\mathcal{T}_k, \mathcal{V}_k)$ *is* $\mathrm{Aut}(G)$*-invariant.*

With a classical understanding of connectivity Theorem 1.1 seems to answer the question (1.4) in the strongest possible way, since the property of being k-inseparable may be considered a minimum requirement for something to be called a $(k+1)$-connected piece. However, the notion of a k-inseparable set ist not suitable to describe the kind of connectivity that is provided by a large grid. For example the graph G depicted in Figure 1.2 has connectivity 2 and contains no k-inseparable set for any $k \geq 4$; but the removal of fewer than 8 vertices will not disconnect *most* pairs of vertices. More precisely, for every set S of fewer than 8 vertices of G, there is one component of $G-S$ which contains more than half of the vertices of G.

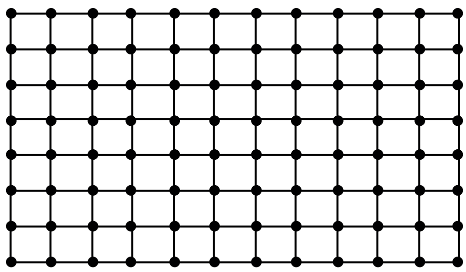

Figure 1.2: A (12×8)-grid, which hosts a tangle of order 8.

Such a kind of connectedness can best be described by the notion of a tangle (of order k) as defined by Robertson and Seymour in [14]. But as we shall see in Chapter 3 and Chapter 5, the notion of a tangle is not suitable to capture k-inseparable sets which are too small in terms of k. This includes, for instance, a maximal clique of size $k+2$, which, as a $(k+1)$-connected subgraph, should clearly be considered a $(k+1)$-connected piece.

We solve this problem by introducing in Chapter 4 the notion of a k-profile, which encompasses both the notion of a tangle of order k and the notion of a $(k-1)$-block. In comparison to other known common generalisations of these two notions (such as the notion of a k-haven), our notion of a k-profile has two major advantages. First, we can define k-profiles both for graphs and matroids, as sets of separations satisfying some simple axioms. In order to do so, we use our theory of separations developed in [4] in a slightly modified version which we present in Chapter 2. And second, (in contrast to k-havens) it is possible to distinguish all the k-profiles of a graph or matroid by a single tree-decomposition of adhesion less than k (see Section 4.4). In particular we are able to show the following result for graphs:

Theorem 1.3. *Every finite graph admits, for every integer k, a canonical tree-decomposition of adhesion at most k that efficiently distinguishes all its distinguishable k-blocks and tangles of order $k+1$, and in which every k-block and every tangle of order $k+1$ inhabits a unique part.*

Within this thesis, a tree-decomposition of a graph (or matroid) is called *canonical* if its construction only depends on the structure of the graph (or matroid). In particular, such a canonical tree-decomposition is invariant under the automorphisms of the graph (or matroid). For tangles alone, we obtain the following strengthening of a result of Robertson and Seymour [14, 10.3]. (They construct a similar tree-decomposition which, however, depends on a fixed vertex enumeration of the graph to 'break ties' between competing separations, and hence is not canonical in our sense.)

Theorem 1.4. *Every finite graph admits a canonical tree-decomposition that efficiently distinguishes all its maximal tangles, and in which every maximal tangle inhabits a unique part.*

Using profiles, we can refine this tree-decomposition further: parts that are inhabited by a unique maximal tangle but still contain more than one (robust) k-block, for some $k \in \mathbb{N}$, can now be split into smaller parts, so that these k-blocks are distinguished too:

Theorem 1.5. *Every finite graph admits a canonical tree-decomposition which, for all k simultaneously, efficiently distinguishes all its distinguishable robust k-blocks and tangles of order $k+1$.*

Extending the Robertson-Seymour theorem cited above, Geelen, Gerards and Whittle [9, 1.1] proved recently that the maximal tangles of a matroid can be distinguished by a single tree-decomposition. Using our theory, we can do the same canonically:

Theorem 1.6. *Every finite matroid has a canonical tree-decomposition that efficiently distinguishes all its maximal tangles, and in which every maximal tangle inhabits a unique part.*

All applications to graphs and matroids, in pariticular the proofs of Theorems 1.3 to 1.6, are treated in Chapter 5.

Chapter 2

Separations and tree-decompositions of finite sets

In this chapter we present the theory of separation systems of [4, Sections 2 to 5]. Our presentation differs slightly from [4] in that our notion of separation is based on an arbitrary finite set V, which need not (but may) be the vertex set of a graph. We have not changed the examples and figures of [4] (which are based on separations of graphs), since we still consider them useful to illuminate the corresponding concepts. In particular, Figures 2.1, 2.2 and 2.3 display the vertices and the edges of a graph.

In Section 2.1 we define the notion of a separation and show how two separations can relate to each other, in particular, we define when a pair of separations is said to be nested. In Section 2.2 we show that every nested set $\mathcal{N}$ of separations of a finite set V can be described by means of a structure tree $\mathcal{T}(\mathcal{N})$. In Section 2.3 we obtain from $\mathcal{T}(\mathcal{N})$ a tree-decomposition of V and show that its parts form $\mathcal{N}$-inseparable sets. In Section 2.4 we study when a given set $\mathcal{I}$ of $\mathcal{S}$-inseparable sets, for some set $\mathcal{S}$ of separations, can be separated by a nested subset $\mathcal{N}$ of $\mathcal{S}$.

2.1 Separations of finite sets

Let V be a finite set. A *separation* of V is an ordered pair (A, B) of subsets $A, B \subseteq V$ such that $A \cup B = V$. A separation (A, B) is *proper* if neither $A \setminus B$ nor $B \setminus A$ is empty. A separation that is not proper is *improper*. The intersection $A \cap B$ is called the *separator* of the separation (A, B). Notice that every separation of a graph (as defined in Chapter 3) is a separation of its vertex set. The converse, however, need not hold.[1]

A separation (A, B) *separates* a set $I \subseteq V$ if I meets both $A \setminus B$ and $B \setminus A$. Two sets I_1, I_2 are *weakly separated* by a separation (A, B) if $I_i \subseteq A$ and

[1]The same is true for matroids: every separation of a matroid M is a separation of its groundset $E(M)$ while the converse need not hold.

$I_{3-i} \subseteq B$ for an $i \in \{1,2\}$. They are *properly separated*, or simply *separated*, by (A,B) if in addition neither I_1 nor I_2 is contained in $A \cap B$.

Given a set $\mathcal{S}$ of separations, a set $I \subseteq V$ is called *$\mathcal{S}$-inseparable* if no separation in $\mathcal{S}$ separates it. A maximal $\mathcal{S}$-inseparable set 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 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 V can relate to each other. The set of all separations of V is partially ordered by

$$(A,B) \leq (C,D) \;:\Leftrightarrow\; A \subseteq C \text{ and } B \supseteq D. \tag{2.1}$$

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:

$$(A,B) \leq (C,D) \;\Leftrightarrow\; (B,A) \geq (D,C). \tag{2.2}$$

Let (C,D) be any proper separation.

No proper separation (A,B) is $\leq$-comparable with both (C,D) and (D,C). In particular, $(C,D) \not\leq (D,C)$. (2.3)

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 (2.2), the other cases all reduce to this case by changing notation: just swap (A,B) with (B,A) or (C,D) or (D,C).

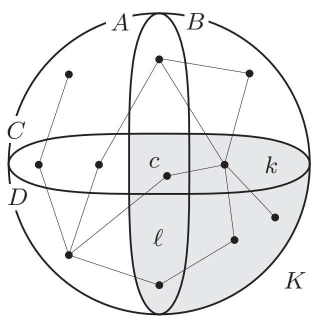

Figure 2.1: The cross-diagram $\{(A,B),(C,D)\}$ with centre c and a corner K and its links k, ℓ.

The way in which two separations relate to each other can be illustrated by a *cross-diagram* as in Figure 2.1. 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, $(A \cup C, B \cap D)$ (in this order) is the corner separation for the corner $B \cap D$ 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:

Any separations (A, B), (C, D) satisfy $(A, B) \leq (A \cup C, B \cap D)$. (2.4)

If (I, J) and (K, L) are distinct corner separations of the same cross-diagram, then $(J, I) \leq (K, L)$. (2.5)

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:

$$(A, B) \leq (C, D) \;\Leftrightarrow\; A \cap D \subseteq B \cap C. \tag{2.6}$$

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.

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 (2.2), this is a symmetrical relation. For example, we saw in (2.4) and (2.5) 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 (2.2). 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 (2.6), 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. (2.7)

In particular:

Neither of two nested separations separates the separator of the other. (2.8)

The converse of (2.8) 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 (K,L) we have either $(K,L) \leq (C,D)$ or $(K,L) \leq (D,C)$.*

Proof. Changing notation as necessary, we may assume that $(A,B) \leq (C,D)$, and that (C,D) is comparable with (E,F).[2] 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 (2.2), $(D,C) \leq (F,E)$. As also $(D,C) \leq (B,A)$, we thus have $E \cup A \subseteq C$ and $F \cap B \supseteq D$ and therfore

$$(J,I) \underset{(2.5)}{\leq} (E \cup A, F \cap B) \leq (C,D)$$

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

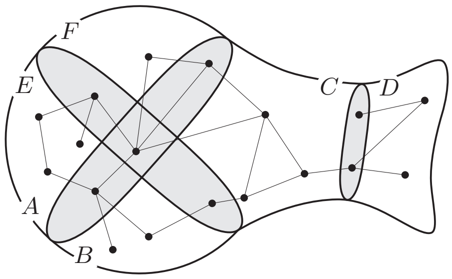

Figure 2.2: Separations as in Lemma 2.2

Figure 2.2 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 2.3.

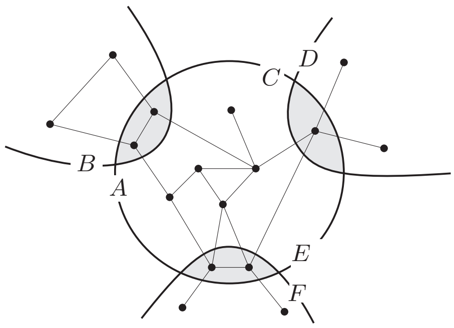

Figure 2.3: Three nested separations not coming from a $\leq$-chain

[2]Note 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.

We need one more lemma.

Lemma 2.3. *Let $\mathcal{N}$ be a set of separations of V 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 $(A, B) \nparallel (C, D)$. By symmetry and Lemma 2.2 we may assume that

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

But then $b' \subseteq F \subseteq A \cap C \subseteq A$, 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 $C \cap D \subseteq b'$. □

2.2 Nested separation systems and tree structure

This section is devoted to the relation of nested sets of separations and trees. First, we show that trees give rise to a nested set of separations.

So consider a tree $\mathcal{T}$ and let $\mathcal{V} := V(\mathcal{T})$. The removal of an edge $e = \{v, w\}$ separates $\mathcal{T}$ into two components. Let $\mathcal{T}_v$ and $\mathcal{T}_w$ denote the component of $\mathcal{T} - e$ containing v and w, respectively. Then each of the two orientations of e gives rise to a separation of $\mathcal{V}$ (in the sense of Section 2.1, these are not separations of $\mathcal{T}$): let $A := V(\mathcal{T}_v)$ and $B := V(\mathcal{T}_w)$, then we associate (v, w) with (A, B) and (w, v) with (B, A). Let $\mathcal{N}(\mathcal{T})$ denote the set of separations that are induced by $\mathcal{T}$. Before we describe the essential properties of this set of separations, we need some more definitions.

We call a set $\mathcal{S}$ of separations *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*.

Proposition 2.4. *Given a tree $\mathcal{T}$, the set $\mathcal{N}(\mathcal{T})$ is a nested separation system.*

Proof. It is obvious by the definition of $\mathcal{N}(\mathcal{T})$ that it is symmetric and that every separation in $\mathcal{N}(\mathcal{T})$ is proper. Thus $\mathcal{N}(\mathcal{T})$ is a separation system. It lasts to show that $\mathcal{N}(\mathcal{T})$ is nested.

So consider distinct elements $(A, B), (C, D) \in \mathcal{N}(\mathcal{T})$. We may assume $(C, D) \neq (B, A)$ since otherwise we clearly have $(A, B) \parallel (C, D)$. Thus (A, B) and (C, D) are induced by (orientations of) distinct edges $e = (a, b)$, $e' = (c, d)$ of $\mathcal{T}$. Then e is contained in either $\mathcal{T}_c$ or $\mathcal{T}_d$, since these are the two components of $\mathcal{T} - e'$. Since nestedness is invariant under flipping a separation we may assume $e \in \mathcal{T}_c$. With the same argument we may assume $e' \in \mathcal{T}_b$ which results in $(A, B) \leq (C, D)$. □

In the rest of this section we will prove the 'converse' of Proposition 2.4. Given an arbitrary set V and a nested system $\mathcal{N}$ of separations of V we aim to describe $\mathcal{N}$ by way of a *structure tree* $\mathcal{T}(\mathcal{N})$, whose oriented edges will correspond to the separations in $\mathcal{N}$. Its set of nodes[3] will correspond to a partition of $\mathcal{N}$. Every group of permutations of V that leaves $\mathcal{N}$ invariant will act on $\mathcal{T}(\mathcal{N})$ as a group of automorphisms. Although our notion of a separation system differs from that of Dunwoody and Krön [7, 6], 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 $\mathcal{T}(\mathcal{N})$ will be to define its nodes. One obvious way to describe the vertices of a tree $\mathcal{T}$ by means of its oriented edges is to identify a vertex v with all oriented edges (v, w) such that $\{v, w\} \in E(\mathcal{T})$. So the idea would be to describe in terms of '$\leq$' as defined in (2.1) when two elements $(A, B), (C, D) \in \mathcal{N}(\mathcal{T})$ correspond to oriented edges that have a common initial vertex. It turns out that this is the case, if and only if we have $(A, B) \sim_{\mathcal{N}} (C, D)$, for $\mathcal{N} = \mathcal{N}(\mathcal{T})$, with

$$(A, B) \sim_{\mathcal{N}} (C, D) \;:\Leftrightarrow \begin{cases} (A, B) = (C, D) \text{ or} \\ (B, A) \text{ is a predecessor of } (C, D) \text{ in } (\mathcal{N}, \leq). \end{cases} \tag{2.9}$$

(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 $x < y < z$.)

For the rest of this section we will fix an arbitrary nested separation system $\mathcal{N}$ and we will write $\sim$ instead of $\sim_{\mathcal{N}}$. We will define the nodes of $\mathcal{T}(\mathcal{N})$ as the equivalence classes with respect to $\sim$.

Lemma 2.5. *The relation $\sim$ is an equivalence relation on $\mathcal{N}$.*

Proof. Reflexivity holds by definition, and symmetry follows from (2.2). 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 (2.2) also

(iii) (D, C) is a predecessor of (A, B);

(iv) (F, E) is a predecessor of (C, D).

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 (2.3) (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), $(X, Y) \not< (C, D)$ and $(D, C) \not< (X, Y)$. Now if $(C, D) \leq (X, Y) < (E, F)$ then by (iv), (C, D) is comparable to both (E, F) and (F, E), contradicting (2.3). 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 (2.3). We have thus shown that (B, A) is a predecessor of (E, F), implying that $(A, B) \sim (E, F)$ as claimed. $\square$

[3]While our graphs have *vertices*, structure trees will have *nodes*.

Note that, by (2.3), the definition of equivalence implies:

Distinct equivalent proper separations are incomparable under $\leq$. (2.10)

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

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

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

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

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,

Every node t of $\mathcal{T}$ is the set of all the separations $(e \cap t)$ such that e is incident with t. In particular, t has degree $|t|$ in $\mathcal{T}$. (2.11)

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

Lemma 2.6. *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 $\sim$. □

And conversely:

Lemma 2.7. *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 $[(E_i,F_i)] = [(F_{i-1},E_{i-1})]$. □

Theorem 2.8. *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 2.6 $(k-1)$ times yields

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

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 set V 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 2.7, $\mathcal{T}$ contains the desired path from $[(A,B)]$ to $[(C,D)]$. □

Corollary 2.9. *If $\mathcal{N}$ is invariant under a group $\Gamma \le \mathrm{Sym}(V)$ of permutations of V, then Γ acts on $\mathcal{T}$ as a group of automorphisms.*

Proof. Any permutation α of V maps separations to separations, and preserves their partial ordering defined in (2.1). 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. □

2.3 From structure trees to tree-decompositions

Throughout this section, $\mathcal{N}$ continues to be an arbitrary nested separation system of a set V. A *tree-decomposition* of V is a pair $(\mathcal{T},\mathcal{V})$ of a tree $\mathcal{T}$ and a family $\mathcal{V} = (V_t)_{t\in\mathcal{T}}$ of subsets $V_t \subseteq V$, one for every node of $\mathcal{T}$, such that:

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

(T2) $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever t_2 lies on the t_1–t_3 path in $\mathcal{T}$.

Recall from Section 2.2 that every oriented edge of a tree induces a separation of its vertex set. Given a tree-decomposition $(\mathcal{T},\mathcal{V})$ of V we can lift a separation of $V(\mathcal{T})$ induced by an oriented edge of the *decomposition tree* $\mathcal{T}$, to a separation of V: if (T_v,T_w) is a separation of $V(\mathcal{T})$ then $\big(\bigcup_{t\in T_v} V_t, \bigcup_{t'\in T_w} V_{t'}\big)$ is a separation of V, due to (T1). We refer to those lifted separations as the separations of V *induced by* the tree-decomposition $(\mathcal{T},\mathcal{V})$.

Our aim now is to show that V has a tree-decomposition with the structure tree $\mathcal{T} = \mathcal{T}(\mathcal{N})$ defined in Section 2.2 as its decomposition tree such that the separations of V that are induced by this tree-decomposition will be precisely the separations in $\mathcal{N}$ identified by those edges in the original definition of $\mathcal{T}$.

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

$$V_t := \bigcap \{\, A \mid (A,B) \in t \,\} \tag{2.12}$$

Example 2.1. Assume that G is a connected graph, and consider as $\mathcal{N}$ the nested set of all proper 1-separations (A,B) and (B,A) of G such that $A \setminus B$ is connected in G. Then $\mathcal{T}$ is very similar to the block-cutvertex tree of G: its

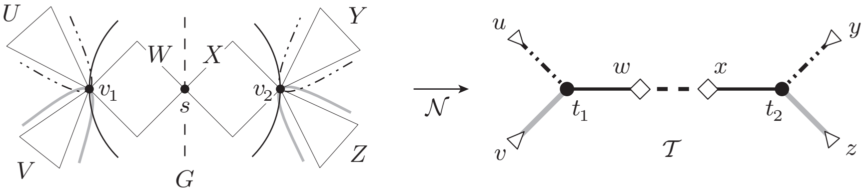

Figure 2.4: $\mathcal{T}$ has an edge for every separation in $\mathcal{N}$. Its nodes correspond to the blocks and some of the cutvertices of G.

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.

In Figure 2.4, 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 $x = [(B, A)]$.

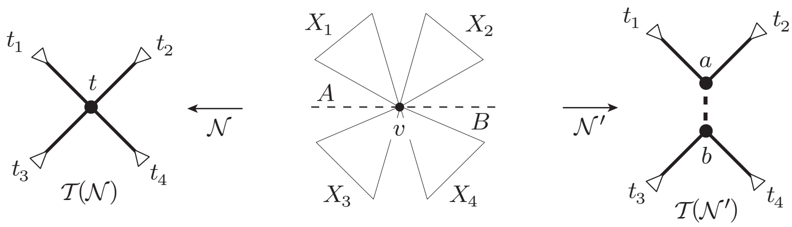

Figure 2.5: $\mathcal{T}' = \mathcal{T}(\mathcal{N}')$ has distinct nodes a, b whose parts in the tree-decomposition $(\mathcal{T}', \mathcal{V})$ coincide: $V_a = \{v\} = V_b$.

In Figure 2.5 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 subsets of V defined in (2.12).

Lemma 2.10. *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 subsets of V), 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 2.1, the $\mathcal{N}$-blocks were the (usual) blocks of G; the hubs were singleton sets consisting of a cutvertex. Example 2.2 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 (2.8), hubs are $\mathcal{N}$-inseparable, so they extend to maximal $\mathcal{N}$-inseparable sets.

Hubs can contain each other properly (Example 2.2 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 2.11. *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 2.12. *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 2.10, V_t is $\mathcal{N}$-inseparable. We show that V_t is maximal in V with this property: that for every element $x \notin V_t$ the set $V_t \cup \{x\}$ is not $\mathcal{N}$-inseparable.

By definition of V_t, any element $x \notin V_t$ lies in $B \setminus A$ for some $(A, B) \in t$. Since t is not a hub node, Lemma 2.11 implies that $V_t \not\subseteq A \cap B$. As $V_t \subseteq A$, this means that V_t has an element in $A \setminus B$. Hence (A, B) separates $V_t \cup \{x\}$, as desired. $\square$

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

Lemma 2.13. *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 (2.8), so in fact $b = A \cap B$ by the maximality of b. We show that $b = V_t$ for $t = [(A, B)]$. By Lemma 2.11, 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

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

(compare the proof of Lemma 2.1). By $(*)$, all corners of the cross-diagram $\{(A, B), (E, F)\}$ other than $B \cap F$ contain elements not in the centre. Hence by (2.7), 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 $(A, B) \le (D, C)$ 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 (2.10) we cannot have $(A, B) \le (X, Y) < (E, F)$, 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 $(X, Y) \le (B, A)$, 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 V, plus some hubs. The following proposition shows what has earned them their name:

Proposition 2.14. *A hub node t has degree at least 3 in $\mathcal{T}$, unless it has the form $t = \{(A, B), (C, D)\}$ with $A \supsetneq D$ 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 $d(t) = |t| \ge 2$; cf. (2.11). 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 $D \subsetneq A$. $\square$

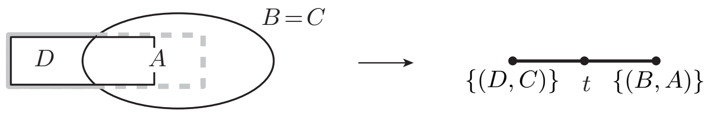

Figure 2.6: A hub node $t = \{(A, B), (C, D)\}$ of degree 2

Figure 2.6 shows that the exceptional situation from Proposition 2.14 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:

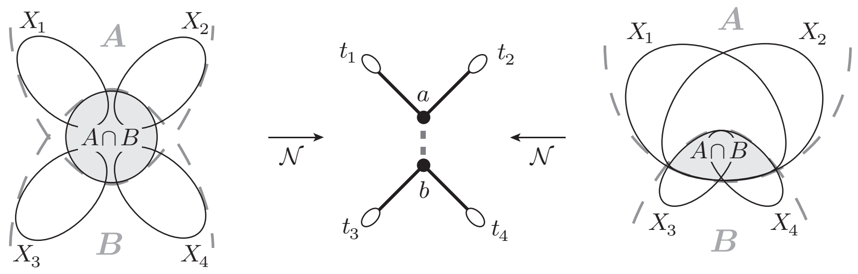

Figure 2.7: The two nested separation systems of Example 2.2, and their common structure tree

Example 2.2. Consider the sets $X_1, \dots, X_4$ shown on the left of Figure 2.7. 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. 2.7, 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 2.15. *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 $V(\mathcal{T}') \setminus \{z\} = V(\mathcal{T}) \setminus \{x, y\}$. [4]

[4] The 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.

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:

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

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 $(C, D) \sim_{\mathcal{N}} (E, F)$ 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 2.7 and Theorem 2.8 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,

$$B \subseteq \bigcap_{(E,F) \in x'} E = V_{x'}. \tag{2.13}$$

By definition of V_x we have $V_x = V_{x'} \cap A$. Hence (2.13) yields $V_{x'} = V_x \cup (B \setminus A)$, and since $A \cap B \subseteq V_x$ by Lemma 2.11, 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,

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

□

We now show that we can lift a separation of $V(\mathcal{T})$ that is induced by an orientation of an edge e of $\mathcal{T}$ to a separation of V in exactly the same way as we did for tree-decompositions in the beginning of this section. This is the separation that (together with its inverse) defined e.

Lemma 2.16. *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 $\mathcal{N}' := \mathcal{N} \smallsetminus e^*$, and let $\mathcal{T}' := \mathcal{T}(\mathcal{N}')$. Then $\mathcal{T}' = \mathcal{T}/e^*$, by Lemma 2.15. 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

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

As $V_z = V_x \cup V_y$ by Lemma 2.15, 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. □

Let us summarize some of our findings from this section. Recall that $\mathcal{N}$ is an arbitrary nested separation system of an arbitrary finite set V. Let $\mathcal{T} := \mathcal{T}(\mathcal{N})$ be the structure tree associated with $\mathcal{N}$ as in Section 2.2, and let $\mathcal{V} := (V_t)_{t\in\mathcal{T}}$ be defined by (2.12).

Theorem 2.17. *The pair $(\mathcal{T}, \mathcal{V})$ is a tree-decomposition of V.*

(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 V induced by the decomposition are precisely those in $\mathcal{N}$.*

(iv) *Every $\mathcal{N}' \subseteq \mathcal{N}$ satsfies $(\mathcal{T}', \mathcal{V}') \preccurlyeq (\mathcal{T}, \mathcal{V})$ for $\mathcal{T}' = \mathcal{T}(\mathcal{N}')$ and $\mathcal{V}' = \mathcal{V}(\mathcal{T}')$.*[5]

Proof. Axiom (T1) follows from Lemma 2.13, because singletons are $\mathcal{N}$-inseparable subsets of V, which extend to $\mathcal{N}$-blocks. For the proof of (T2), 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 2.16 and 2.11 imply that $V_{t_1} \cap V_{t_3} \subseteq A \cap B \subseteq V_{t_2}$.

Statement (i) is Lemma 2.13. Assertion (ii) is Lemma 2.12. Assertion (iii) follows from Lemma 2.16 and the definition of the edges of $\mathcal{T}$. Statement (iv) follows by repeated application of Lemma 2.15. □

2.4 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 subsets of V. As we saw in Sections 2.2 and 2.3, such a nested

[5]See the Introduction for the definition of $(\mathcal{T}', \mathcal{V}') \preccurlyeq (\mathcal{T}, \mathcal{V})$.

subsystem will then define a tree-decomposition of V, and the sets from $\mathcal{I}$ will come to lie in different parts of that decomposition.

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 V be the vertex set of 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 $I_2 \subseteq F \subseteq B \cap D$ and $E \supseteq A \cup C$.*

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

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 $(C,D) \leq (E,F)$ or $(D,C) \leq (E,F)$. In particular, extremal separations are nested with all other separations in $\mathcal{S}$. Being extremal implies being $\leq$-maximal in $\mathcal{S}$; if $\mathcal{S}$ is nested, extremality and $\leq$-maximality are equivalent. If $(E,F) \in \mathcal{S}$ is extremal, then F 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 V 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 V corresponding to edges of $\mathcal{T}$.[6]

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.*

[6] More 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 B 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 2.18. *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 $\leq$-maximal $(A,B) \in \mathcal{R}$ is extremal in $\mathcal{R}$. In particular, if $\mathcal{R} \neq \emptyset$ then $\mathcal{R}$ contains an extremal separation.*

Proof. Consider a $\leq$-maximal 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 maximality of (A,B) implies that $(C,D) \leq (A,B)$ or $(D,C) \leq (A,B)$, 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 $I_2 \subseteq F \subseteq B \cap D$ and $E \supseteq A \cup C$, and hence $(A,B) \leq (E,F)$ as well as $(C,D) \leq (E,F)$. By the maximality of (A,B), this yields $(C,D) \leq (E,F) = (A,B)$ 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 2.19. *Let $\mathcal{S}$ be any separation system that separates some set $\mathcal{I}$ of $\mathcal{S}$-inseparable sets 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

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

By Lemma 2.18 we have $\overline{\mathcal{E}} \neq \emptyset$, and by definition of extremality all separations in $\overline{\mathcal{E}}$ are nested with all separations in $\mathcal{R}$. In particular, $\overline{\mathcal{E}}$ is nested.

Let

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

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 $I_2 \in \mathcal{I}_{\mathcal{E}}$. Pick $(A,B) \in \mathcal{R}$ with $I_1 \subseteq A$ and $I_2 \subseteq B$. Since $I_2 \in \mathcal{I}_{\mathcal{E}}$, there is an $(E,F) \in \mathcal{E}$ such that $I_2 \subseteq F$. By the extremality of (E,F) we have either $(A,B) \leq (E,F)$, in which case $I_1 \subseteq A \subseteq E$ and $I_2 \subseteq F$, or we have $(B,A) \leq (E,F)$, in which case $I_2 \subseteq B \cap F \subseteq E \cap F$. 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 2.19 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 permutation of V that leave $\mathcal{S}$ and $\mathcal{I}$ invariant.

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

Corollary 2.20. *Let $\mathcal{S}$ and $\mathcal{I}$ be as in Theorem 2.19, and let $\mathcal{N}(\mathcal{S}, \mathcal{I})$ be the nested subsystem of $\mathcal{S}$ constructed in the proof. Then for every permutation α of V 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 permutations of V, 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 $\alpha \in \Gamma$. □

Chapter 3

The k-blocks of graphs

The essential parts of this chapter are based on [3] and [4]. In Section 3.1 we give alternative[1] definitions of separations and tree-decompositions of graphs, using the language of Chapter 2, and introduce some further related terminology that we need in subsequent sections. In Section 3.2 we present the examples of [3, Section 3] that demonstrate three different types of k-blocks. In Section 3.3 we come back to the main subject of this thesis: based on [4, Section 6] we present how to distinguish the k-blocks of a graph by means of a tree-decomposition. In Sections 3.4 and 3.5 we show how to force the existence of a k-block in a graph by means of minimum and average degree (see [3, Sections 4 and 5]).

3.1 Separations and tree-decompositions of graphs

Recall from the Introduction that a graph is a pair $G = (V, E)$ of a set V of vertices and a set E of edges. In the language of Chapter 2, a *separation of* G is a separation of its set of vertices V which does not separate any edge of G. That is, a separation (A, B) of V is a separation of G, if and only if $\{v, w\} \subseteq A$ or $\{v, w\} \subseteq B$, for every edge $\{v, w\} \in E$. A separation (A, B) of G is *tight* if every vertex $v \in A \cap B$ in its separator has neighbors both in $A \setminus B$ and $B \setminus A$. A set of separations is *tight* if all of its elements are tight.

The *order* of a separation of G is the size of its separator. A separation of order k is called a k-*separation*. A simple calculation yields the following:

Lemma 3.1. *Given any two separations (A, B) and (C, D) of G, the orders of the separations $(A \cap C, B \cup D)$ and $(B \cap D, A \cup C)$ sum to $|A \cap B| + |C \cap D|$.* $\square$

A set I of at least $k+1$ vertices is k-*inseparable*, if for every separation (A, B) of G of order at most k we have $I \subseteq A$ or $I \subseteq B$. A maximal k-inseparable set of vertices is called a k-*block*. Now let $\mathcal{S}_k$ denote the set of all separations of G of order at most k. Then a set I of vertices is k-inseparable if and only if it is both $\mathcal{S}_k$-inseparable and large with respect to $\mathcal{S}_k$, and I is a k-block, if and only if it is a large $\mathcal{S}_k$-block.

[1] but equivalent to those given in the Introduction

A *tree-decomposition of* G is a tree-decomposition $(\mathcal{T}, \mathcal{V})$ of V which in addition to the two axioms (T1) and (T2) given on Page 14, satisfies the following axiom (T3):

(T3) for every edge $e \in G$ there exists a $t \in \mathcal{T}$ such that both ends of e lie in V_t.

Now consider a tree-decomposition $(\mathcal{T}, \mathcal{V})$ of V and let $\mathcal{N}$ be the (nested) set of separations of V induced by $(\mathcal{T}, \mathcal{V})$. Then by Lemma 2.10 and Thorem 2.17, the pair $(\mathcal{T}, \mathcal{V})$ is a tree-decomposition of G, if and only if every edge $e \in G$ is $\mathcal{N}$-inseparable, which in turn is the case if and only if every separation in $\mathcal{N}$ is a separation of G. That is we have the following:

A tree-decomposition of V that is obtained from a nested system $\mathcal{N}$ of separations of G is a tree-decomposition of G. (3.1)

The set of separations of V that is induced by a tree-decomposition of G is a set of separations of G. (3.2)

The *width* of a tree-decomposition $(\mathcal{T}, \mathcal{V})$ is the number $\max_{t \in \mathcal{T}} |V_t| - 1$, and the *tree-width* of G is the least width of any tree-decomposition of G.

The intersections $V_t \cap V_{t'}$ of 'adjacent' parts in a tree-decomposition $(\mathcal{T}, \mathcal{V})$ (those for which tt' is an edge of $\mathcal{T}$) are its *adhesion sets*; the maximum size of such a set is the *adhesion* of $(\mathcal{T}, \mathcal{V})$. The *interior* of a part V_t, denoted by $\mathring{V}_t$, is the set of those vertices in V_t that lie in no adhesion set. By (T2), we have $\mathring{V}_t = V_t \setminus \bigcup_{t' \neq t} V_{t'}$. It is easy to see that we could have alternaitively defined the adhesion of a tree-decomposition by its induced set of separations.

Lemma 3.2. *The adhesion of a tree-decomposition is equal to the maximum order of a separation it induces.* □

A tree-decomposition $(\mathcal{T}, \mathcal{V})$ of a graph G is *lean* if for any nodes $t_1, t_2 \in \mathcal{T}$, not necessarily distinct, and vertex sets $Z_1 \subseteq V_{t_1}$ and $Z_2 \subseteq V_{t_2}$ such that $|Z_1| = |Z_2| =: \ell$, either G contains ℓ disjoint Z_1–Z_2 paths or there exists an edge $tt' \in t_1 \mathcal{T} t_2$ with $|V_t \cap V_{t'}| < \ell$. Since there is no such edge when $t_1 = t_2 =: t$, this implies in particular that, for every part V_t, any two subsets $Z_1, Z_2 \subseteq V_t$ of some equal size ℓ are linked in G by ℓ disjoint paths. (However, the parts need not be ℓ-inseparable for any large ℓ; see Section 3.2.)

We call a tree-decomposition $(\mathcal{T}, \mathcal{V})$ *k-lean* if none of its parts contains another, it has adhesion at most k, and for any nodes $t_1, t_2 \in \mathcal{T}$, not necessarily distinct, and vertex sets $Z_1 \subseteq V_{t_1}$ and $Z_2 \subseteq V_{t_2}$ such that $|Z_1| = |Z_2| =: \ell \leq k + 1$, either G contains ℓ disjoint Z_1–Z_2 paths or there exists an edge $tt' \in t_1 \mathcal{T} t_2$ with $|V_t \cap V_{t'}| < \ell$.

Thomas [16] proved that every graph G has a lean tree-decomposition whose width is no greater than the tree-width of G. By considering only separations of order at most k one can adapt the short proof of Thomas's theorem given by Bellenbaum and Diestel in [1] to yield the following:

Theorem 3.3. *Every graph has a k-lean tree-decomposition.* □

3.2 Examples of k-blocks

In this section we discuss three different types of k-block (see [3, Section 3]).

Example 3.1. The vertex set of any $(k{+}1)$-connected subgraph is k-inseparable, and hence contained in a k-block.

While a k-block as in Example 3.1 derives much or all of its inseparability from its own connectivity as a subgraph, the k-block in our next example will form an independent set. It will derive its inseparability from the ambient graph, a large grid to which it is attached.

Example 3.2. Let $k \geq 4$, and let H be a large $(m \times n)$-grid, with $m, n \geq (k+1)^2$ say. Let G be obtained from H by adding a set $X = \{x_1, \ldots, x_{k+1}\}$ of new vertices, joining each x_i to at least $k+1$ vertices on the grid boundary that form a (horizontal or vertical) path in H so that every grid vertex obtains degree 4 in G (Figure 3.1). We claim that X is a k-block of G, and is its only k-block.

Any grid vertex can lie in a common k-block of G only with its neigbhours, because these separate it from all the other vertices. As any k-block has at least $k+1 \geq 5$ vertices but among the four G-neighbours of a grid vertex at least two are non-adjacent grid vertices, this implies that no k-block of G contains a grid vertex. On the other hand, every two vertices of X are linked by $k+1$ independent paths in G, and hence cannot be separated by at most k vertices. Hence X is k-inseparable, maximally so, and is thus the only k-block of G.

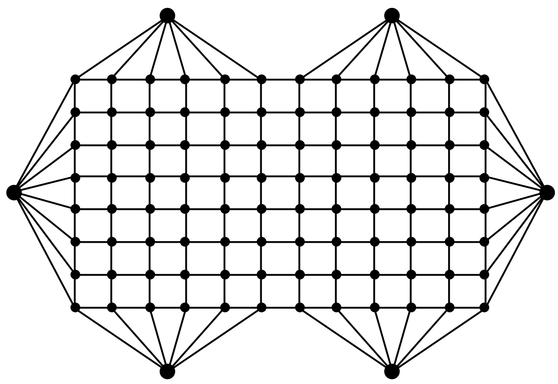

Figure 3.1: The six outer vertices form a 5-block.

In the discussion of Example 3.2 we saw that none of the grid vertices lies in a k-block. In particular, the grid itself has no k-block when $k \geq 4$. Since every two inner vertices of the grid, those of degree 4, are joined in the grid by 4 independent paths, they form a 3-inseparable set (which is clearly maximal):

Example 3.3. The inner vertices of any large grid H form a 3-block in H. However, H has no k-block for any $k \geq 4$.

The k-block defined in Example 3.2 gives rise to a tangle of large order (see Section 5.1), the same as the tangle specified by the grid H. This is in contrast to our last two examples, where the inseparability of the k-block will again lie in the ambient graph but in a way that need not give rise to a non-trivial tangle. (See Section 5.1 for when it does.) Instead, the paths supplying the required connectivity will live in many different components of the subgraph into which the k-block splits the original graph.

Example 3.4. Let X be a set of $n \geq k+1$ isolated vertices. Join every two vertices of X by many (e.g., $k+1$) independent paths, making all these internally disjoint. Then X will be a k-block in the resulting graph.

Example 3.4 differs from Example 3.2 in that its graph has a tree-decomposition whose only part of order ≥ 3 is X. Unlike the grid in Example 3.2, the paths providing X with its external connectivity do not between them form a subgraph that is in any sense highly connected. We can generalize this as follows:

Example 3.5. Given $n \geq k+1$, consider a tree T in which every non-leaf node has $\binom{n}{k}$ successors. Replace each node t by a set V_t of n isolated vertices. Whenever t' is a successor of a node t in T, join $V_{t'}$ to a k-subset $S_{t'}$ of V_t by k independent edges, so that these $S_{t'}$ are distinct sets for different successors t' of t. For every leaf t of T, add edges on V_t to make it complete. The k-blocks of the resulting graph G are all the sets V_t $(t \in T)$, but only the sets V_t with t a leaf of T induce any edges.

3.3 Separating the k-blocks of a graph

As in [4] we now apply the theory developed in Chapter 2 to our original problem, of how to 'decompose a graph G into its $(k+1)$-connected components'. In the language of Section 2.4, 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 2.4 rest on the assumption that the set $\mathcal{R}$ of $\mathcal{I}$-relevant separations in $\mathcal{S}$ separates $\mathcal{I}$ well (Lemma 2.18). 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 $b_2 \subseteq F \subseteq B \cap D$?

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 $|A \cap B| + |C \cap D| = 2k$, by Lemma 3.1, so they are both exactly k. We can thus take as (E, F) the corner separation $(A \cup C, B \cap D)$.

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 2.3[2]. 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 $\mathrm{Aut}(G)$ of G.

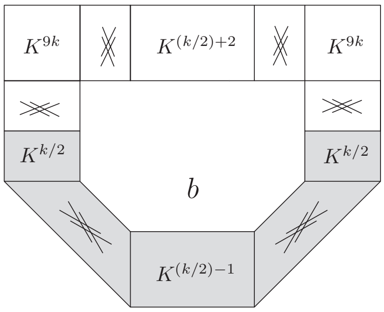

Figure 3.2: A horizontal k-separation needed to distinguish two k-blocks, crossed by a vertical $(k+1)$-separation needed to distinguish two $(k+1)$-blocks.

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 3.2 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 3.2 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 3.2, inside the union

[2]These are tree-decomposition *of* V due to Theorem 2.17 but also *of* G by (3.1).

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*[3] if for every k-separation (C,D) with $U \subseteq D$ and every separation (A,B) of order at most K such that $(A,B) \not\parallel (C,D)$ and

$$|\partial(A \cap D)| < k > |\partial(B \cap D)|\,, \tag{3.3}$$

we have either $U \subseteq A$ or $U \subseteq B$. By $U \subseteq D$ and (3.3), 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. 3.3): exactly the situation of b in Figure 3.2.

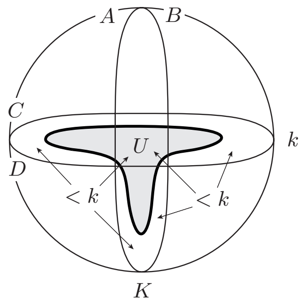

Figure 3.3: The shaded set U is k-inseparable but not K-robust.

It is obvious from the definition of robustness that

$$\textit{for } k \geq K, \textit{ every } k\textit{-inseparable set is } K\textit{-robust.} \tag{3.4}$$

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 3.4. *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 3.3 has size at most $\frac{3}{2}k - 1$, regardless of the order of (A,B). Let $\ell := |(A \cap B) \smallsetminus C|$ be the size of the common link of the corners $A \cap D$ and $B \cap D$. By $|C \cap D| = k$ and (3.3) we have $2\ell \leq k - 2$, so $|T| = k + \ell \leq \frac{3}{2}k - 1$ as desired. □

For the remainder of this section, 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:

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

Lemma 3.5. *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.[4] 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 3.6. *For every finite graph G and every integer $k \geq 0$ there is a tight, nested, and* $\mathrm{Aut}(G)$*-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 $\mathrm{Aut}(G)$-invariant;

[4] Unlike in the definition just before Theorem 2.19, we no longer require that the blocks we wish to separate be $\mathcal{S}$-inseparable for the entire set $\mathcal{S}$.

(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 3.5 (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 $k' = 0, \ldots, k-1$. (For $k = 0$ we have nothing but the definiton of $\mathcal{N}_{-1} := \emptyset$, which has $V(G)$ as its unique $\mathcal{N}_{-1}$-block.) Let us show the following:

$$\text{For all } 0 \leq \ell \leq k \text{, any two } K\text{-robust } \ell\text{-blocks } b_1, b_2 \text{ that are not distinguished by } \mathcal{N}_{\ell-1} \text{ satisfy } \kappa(b_1, b_2) = \ell. \tag{3.5}$$

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 3.5 (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 (3.5) for $\ell = k$, the separations in $\mathcal{S}_b$ are tight. Our aim is to apply Theorem 2.19 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 2.19, let us prove that it will be useful: that the nested separation system $\mathcal{N}_b \subseteq \mathcal{S}_b$ it yields can distinguish[5] 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

[5] As the elements of $\mathcal{I}_b$ are k-blocks, we have two notions of 'distinguish' that could apply: the definition given before Theorem 2.19, or that given before Theorem 3.6. 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.

$I_2 \subseteq B$. 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 (3.5). 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 $(C, D) \in \mathcal{N}_\ell \smallsetminus \mathcal{N}_{\ell-1}$. 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 (3.5),

$$\kappa(J_1, J_2) = \ell. \tag{3.6}$$

Let us show that we may assume the following:

$$\text{The corner separations of the corners } A \cap C \text{ and } B \cap C \text{ are nested with every separation } (C', D') \in \mathcal{N}_{k-1} \text{ that } (A, B) \text{ is nested with.} \tag{3.7}$$

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 (3.7).

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), (3.7) 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 (3.6), completing the proof of $(*)$.

Let us now verify the premise of Theorem 2.19:

$(**)$ *$\mathcal{S}_b$ 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 $(A \cup C, B \cap D) \in \mathcal{S}_b$.

By (3.5) 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 $(A \cup C, B \cap D)$ 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 2.19 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 $\mathrm{Aut}(G)$-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 2.20 yields $(\mathcal{N}_b)^\alpha = \mathcal{N}_{b^\alpha}$. Thus, $\mathcal{N}_{\mathcal{B}}$ is $\mathrm{Aut}(G)$-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.[6] For $k = |G|$, Theorem 3.6 then yields our 'unified' nested separation system that separates all robust blocks by a separation of the lowest possible order:

Corollary 3.7. *For every finite graph G there is a tight, nested, and* $\mathrm{Aut}(G)$*-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 3.6 and its proof into tree-decompositions:

Theorem 3.8. *For every finite graph G and every integer K there is a sequence $(\mathcal{T}_k, \mathcal{V}_k)_{k \le K}$ of tree-decompositions such that, for all $k \le K$,*

(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 $(\mathcal{T}_{k-1}, \mathcal{V}_{k-1}) \preccurlyeq (\mathcal{T}_k, \mathcal{V}_k)$;*

(v) $\mathrm{Aut}(G)$ *acts on $\mathcal{T}_k$ as a group of automorphisms.*

Proof. Consider the nested separation system $\mathcal{N}_K$ given by Theorem 3.6. 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 3.6, $\mathcal{N}_K$ is $\mathrm{Aut}(G)$-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 of V associated with $\mathcal{N}_k$ as in Section 2.3. By (3.1) this is also a tree-decomposition of G. Then (v) holds by Corollary 2.9, (iii) and (iv) by Theorem 2.17 (iii) and (iv). By (iii) and [5, 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 3.6, $\mathcal{N}_k$ distinguishes every two K-robust k-blocks, which implies (ii) by (i) and Theorem 2.17 (iii). $\square$

[6]By Lemma 3.5 (i), this is equivalent to saying that they are robust $r(b_i)$-blocks, that is, K-robust $r(b_i)$-blocks for $K = |G|$.

From Theorem 3.8 we can finally deduce two of the main results announced in the Introduction, Theorems 1.1 and 1.2.

Theorem 1.1 follows by taking as K the integer k given in Theorem 1.1, and then considering the decomposition $(\mathcal{T}_k, \mathcal{V}_k)$ for $k = K$. Indeed, consider two k-blocks b_1, b_2 that Theorem 1.1 claims are distinguished efficiently by $(\mathcal{T}_k, \mathcal{V}_k)$. By Theorem 3.8 (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 3.5 (i), the k'-blocks $b'_1 \supseteq b_1$ and $b'_2 \supseteq b_2$ are again K-robust. Hence by Theorem 3.8 (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 3.8 (iii), this set has size at most k', and by Theorem 3.8 (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 1.2 follows from Theorem 3.8 for $K = |G|$; recall that robust k-blocks are K-robust for $K = |G|$.

3.4 Forcing k-blocks by minimum degree conditions

Throughout this section, let $G = (V, E)$ be a fixed non-empty graph, and $k \geq 0$ an integer. We ask what minimum degree will force G to contain a k-block.

Without any further assumptions on G we shall see that $\delta(G) \geq 2k$ will be enough. If we assume that G is k-connected – an interesting case, since for such G the parameter k is minimal such that looking for k-blocks can be non-trivial – we find that $\delta(G) > \frac{3}{2}k - 1$ suffices. If G is k-connected but contains no triangle, even $\delta(G) \geq k + 1$ will be enough. Note that this is best possible in the (weak) sense that the vertices in any k-block will have to have degree at least $k + 1$, except in some very special cases that are easy to describe.

Conversely, we construct a k-connected graph of minimum degree $\lfloor \frac{3}{2}k - 1 \rfloor$ that has no k-block. So our second result above is sharp.

We shall often use the fact that a vertex of G together with k or more of its neighbours forms a k-inseparable set as soon as these neighbours are pairwise not separated by $\leq k$ vertices. Let us state this as a lemma:

Lemma 3.9. *Let $v \in V$ and $N \subseteq N(v)$ with $|N| \geq k$. If no two vertices of N are separated in G by at most k vertices, then $N \cup \{v\}$ lies in a k-block.* □

Here, then, is our first sufficient condition for the existence of a k-block:

Theorem 3.10. *If $\delta(G) \geq 2k$, then G has a k-block. This can be chosen to be connected in G and of size at least $\delta(G) + 1 - k$.*

Proof. If $k = 0$, then the assertion follows directly. So we assume $k > 0$. By Theorem 3.3, G has a k-lean tree-decomposition $(\mathcal{T}, \mathcal{V})$, say with $\mathcal{V} = (V_t)_{t \in \mathcal{T}}$. Pick a leaf t of $\mathcal{T}$. (If $\mathcal{T}$ has only one node, we count it as a leaf.) Write $A_t := V_t \cap \bigcup_{t' \neq t} V_{t'}$ for the attachment set of V_t. As V_t is not contained in any other part of $(\mathcal{T}, \mathcal{V})$, we have $\mathring{V}_t = V_t \setminus A_t \neq \emptyset$ by (T3). By our degree assumption and $|A_t| \leq k$, every vertex in $\mathring{V}_t$ has k neighbours in $\mathring{V}_t$. Thus, $|\mathring{V}_t| \geq k + 1 \geq 2$.

We prove that $\mathring{V}_t$ extends to a k-block $B \subseteq V_t$ that is connected in G. Pick distinct vertices $v, v' \in \mathring{V}_t$. Let N be a set of k neighbours of v, and N' a set of k neighbours of v'. Note that $N \cup N' \subseteq V_t$. As our tree-decomposition is k-lean, there are $k+1$ disjoint paths in G between the $(k+1)$-sets $N \cup \{v\}$ and $N' \cup \{v'\}$. Hence v and v' cannot be separated in G by at most k other vertices.

We have thus shown that $\mathring{V}_t$ is k-inseparable. In particular, A_t does not separate it, so $\mathring{V}_t$ is connected in G. Let B be a k-block containing $\mathring{V}_t$. As A_t separates $\mathring{V}_t$ from $G \setminus V_t$, we have $B \subseteq V_t$. Every vertex of B in A_t sends an edge to $\mathring{V}_t$, since otherwise the other vertices of A_t would separate it from $\mathring{V}_t$. Hence B is connected. Since every vertex in $\mathring{V}_t$ has at least $\delta(G) - k$ neighbours in $\mathring{V}_t \subseteq B$, we have the desired bound of $|B| \geq \delta(G) + 1 - k$. □

We do not know whether the degree bound of $\delta(G) \geq 2k$ in Theorem 3.10 is sharp. The largest minimum degree known of a graph without a k-block is $\lfloor \frac{3}{2}k - 1 \rfloor$. This graph (Example 3.6 below) is k-connected, and we shall see that k-connected graphs of larger minimum degree do have k-blocks (Theorem 3.14). Whether or not graphs of minimum degree between $\frac{3}{2}k - 1$ and $2k$ and connectivity $< k$ must have k-blocks is unknown to us.

It is also conceivable that the smallest minimum degree that will force a *connected* k-block – or at least a connected k-inseparable set, as found by our proof of Theorem 3.10 – is indeed $2k$ but possibly disconnected k-blocks can be forced by a smaller value of δ.

The degree bound of Theorem 3.10 can be reduced by imposing additional conditions on G. Our next aim is to derive a better bound on the assumption that G is k-connected, for which we need a few lemmas.

We say that a k-separation (A, B) is *T-shaped* (Fig. 3.4) if it is a proper separation and there exists another proper k-separation (C, D) such that $A \setminus B \subseteq C \cap D$ as well as $|A \cap C| \leq k$ and $|A \cap D| \leq k$. Obviously, (A, B) is T-shaped witnessed by (C, D) if and only if the two separations $(A \cap C, B \cup D)$ and $(A \cap D, B \cup C)$ have order at most k and are improper separations.

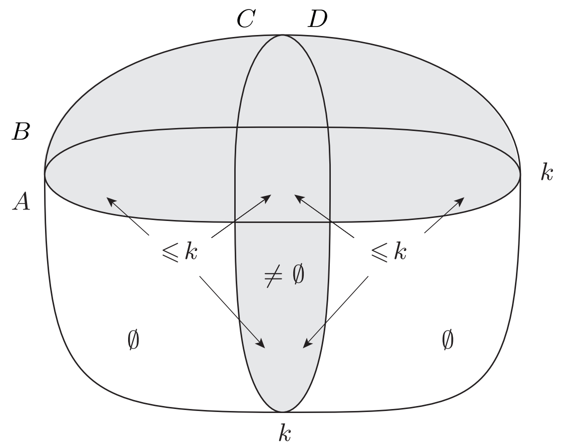

Figure 3.4: The separation (A, B) is T-shaped

Lemma 3.11. *If (A, B) is a T-shaped k-separation in G, then $|A| \leq \frac{3}{2}k$.*

Proof. Let (C, D) witness that (A, B) is T-shaped. Then

$$|A| \leq |A \cap B| + |(C \cap D) \setminus B| \leq k + \tfrac{1}{2}(2k - k) = \tfrac{3}{2}k.$$ □

When a k-separation (A, B) is T-shaped, no k-block of G can lie in A: with (C, D) as above, it would have to lie in either $A \cap C$ or $A \cap D$, but both these are too small to contain a k-block. Conversely, one may ask whether every proper k-separation (A, B) in a k-connected graph such that A contains no k-block must be T-shaped, or at least give rise to a T-shaped k-separation (A', B') with $A' \subseteq A$. However, this is not the case, as can be shown by counterexample.

Interestingly, though, a global version of this does hold: a T-shaped k-separation must occur *somewhere* in every k-connected graph that has no k-block. More precisely, we have the following:

Lemma 3.12. *If G is k-connected, the following statements are equivalent:*

(i) *every proper k-separation of G separates two k-blocks;*

(ii) *no k-separation of G is T-shaped.*

Proof. We first assume (i) and show (ii). If (ii) fails, then G has a k-separation (A, B) that is T-shaped, witnessed by (C, D) say. We shall derive a contradiction to (i) by showing that A contains no k-block. If A contains a k-block, it lies in either $A \cap C$ or $A \cap D$, since no two of its vertices are separated by (C, D). By the definition of T-shaped, none of these two cases can occur, a contradiction.

Let us now assume (ii) and show (i). If (i) fails, there is a proper k-separation (A, B) such that A contains no k-block. Pick such an (A, B) with $|A|$ minimum. Since (A, B) is proper, there is a vertex $v \in A \setminus B$. Since G is k-connected, v has at least k neighbours, all of which lie in A. As A contains no k-block, Lemma 3.9 implies that there is a proper k-separation (C, D) that separates two of these neighbours. Then v must lie in $C \cap D$.

We first show that either $(A \cap C, B \cup D)$ has order at most k and $(A \cap C) \setminus (B \cup D) = \emptyset$ or $(B \cap D, A \cup C)$ has order at most k and $(B \cap D) \setminus (A \cup C) = \emptyset$. Let us assume that the first of these fails; then either $(A \cap C, B \cup D)$ has order $> k$ or $(A \cap C) \setminus (B \cup D) \neq \emptyset$. In fact, if the latter holds then so does the former: otherwise $(A \cap C, B \cup D)$ is a proper k-separation that contradicts the minimality of $|A|$ in the choice of (A, B). (We have $|A \cap C| < |A|$, since v has a neighbour in $A \setminus C$.) Thus, $(A \cap C, B \cup D)$ has order $> k$. As $|A \cap B| + |C \cap D| = 2k$, this implies by Lemma 3.1 that the order of $(B \cap D, A \cup C)$ is strictly less than k. As G is k-connected, this means that $(B \cap D, A \cup C)$ is not a proper separation, i.e., that $(B \cap D) \setminus (A \cup C) = \emptyset$ as claimed.

By symmetry, we also get the analogous statement for the two separations $(A \cap D, B \cup C)$ and $(B \cap C, A \cup D)$. But this means that one of the separations (A, B), (B, A), (C, D) and (D, C) is T-shaped, contradicting (ii). $\square$

Our next lemma says something about the size of the k-blocks we shall find.

Lemma 3.13. *If G is k-connected and $|A| > \frac{3}{2}k$ for every proper k-separation (A, B) of G, then either V is a k-block or G has two k-blocks of size at least $\mu := \min\{\, |A| : (A, B)$ is a proper k-separation$\}$ that are connected in G.*

Proof. By assumption and Lemma 3.11, G has no T-shaped k-separation, so by Lemma 3.12 every side of a proper k-separation contains a k-block.

By Theorem 3.3, G has a k-lean tree-decomposition $(\mathcal{T}, \mathcal{V})$, with $\mathcal{V} = (V_t)_{t \in \mathcal{T}}$ say. Unless V is a k-block, in which case we are done, this decomposition has at least two parts: since there exist two $(k+1)$-sets in V that are separated by some k-separation, the trivial tree-decomposition with just one part would not be k-lean.

So $\mathcal{T}$ has at least two leaves, and for every leaf t the separation $(A,B) := \big(V_t, \bigcup_{t'\neq t} V_{t'}\big)$ is a proper k-separation. It thus suffices to show that $A = V_t$ is a k-block; it will clearly be connected (as in the proof of Theorem 3.10).

As remarked at the start of the proof, there exists a k-block $X \subseteq A$. If $X \neq A$, then A has two vertices that are separated by a k-separation (C,D); we may assume that $X \subseteq C$, so $X \subseteq A \cap C$.

If $(A\cap C, B\cup D)$ has order $\leq k$, it is a proper separation (as $X \subseteq A\cap C$ has size $> k$); then its separator S has size exactly k, since G is k-connected. By the choice of (C,D) there is a vertex v in $(D \setminus C) \cap A$. The $k+1$ vertices of $S \cup \{v\} \subseteq A$ are thus separated in G by the k-set $C \cap D$ from $k+1$ vertices in $X \subseteq A \cap C$, which contradicts the leanness of $(\mathcal{T}, \mathcal{V})$ for $V_t = A$.

So the order of $(A\cap C, B\cup D)$ is at least $k+1$. By Lemma 3.1, the order of $(B\cap D, A\cup C)$ must then be less than k, so by the k-connectedness of G there is no k-block in $B \cap D$.

The k-block X' which D contains (see earlier) thus lies in $D \cap A$. So A contains two k-blocks X and X', and hence two vertex sets of size $k+1$, that are separated by (C,D), which contradicts the k-leanness of $(\mathcal{T},\mathcal{V})$. $\square$

Theorem 3.14. *If G is k-connected and $\delta(G) > \frac{3}{2}k - 1$, then either V is a k-block or G has at least two k-blocks. These can be chosen to be connected in G and of size at least $\delta(G)+1$.*

Proof. For every proper k-separation (A,B) we have a vertex of degree $> \frac{3}{2}k-1$ in $A \setminus B$, and hence $|A| \geq \delta(G)+1 > \frac{3}{2}k$. The assertion now follows from Lemma 3.13. $\square$

To show that the degree bound in Theorem 3.14 is sharp, let us construct a k-connected graph H with $\delta(H) = \lfloor \frac{3}{2}k - 1 \rfloor$ that has no k-block.

Example 3.6. Let H_n be the ladder that is a union of $n \geq 2$ squares (formally: the cartesian product of a path of length n with a K^2).

For even k, let H be the lexicographic product of H_n and a complete graph $K = K^{k/2}$, i.e., the graph with vertex set $V(H_n) \times V(K)$ and edge set

$$\{\, (h_1,x)(h_2,y) \mid \text{either } h_1 = h_2 \text{ and } xy \in E(K) \text{ or } h_1h_2 \in E(H_n) \,\},$$

see Figure 3.5. This graph H is k-connected and has minimum degree $\frac{3}{2}k-1$. But it contains no k-block: among any $k+1$ vertices we can find two that are separated in H by a k-set of the form $V_{h_1} \cup V_{h_2}$, where $V_h := \{(h,x) \mid x \in K\}$.

If k is odd, let H' be the graph H constructed above for $k-1$, and let H be obtained from H' by adding a new vertex and joining it to every vertex of H'. Clearly, H is again k-connected and has minimum degree $\lfloor \frac{3}{2}k-1 \rfloor$, and it has no k-block since H has no $(k-1)$-block.

Our next example shows that the connectivity bound in Theorem 3.14 is sharp: we construct for every odd k a $(k-1)$-connected graph H of minimum degree $\lfloor \frac{3}{2}k \rfloor$ whose largest k-blocks have size $k+1 < \delta(H)+1$.

Example 3.7. Let H_n be as in Example 3.6. Let H be obtained from H_n by replacing the degree-two vertices of H_n by complete graphs of order $(k+1)/2$

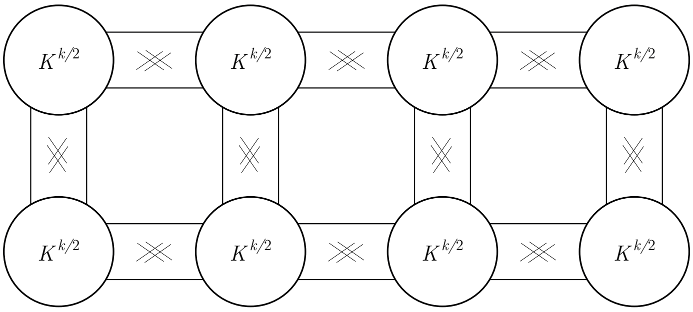

Figure 3.5: A k-connected graph without a k-block.

and its degree-three vertices by complete graphs of order $(k-1)/2$, joining vertices of different complete graphs whenever the corresponding vertices of H_n are adjacent. The minimum degree of this graph is $\lfloor \frac{3}{2}k \rfloor$, but it has only two k-blocks: the two K^{k+1}s at the extremes of the ladder.

We do not know whether the assumption of k-connectedness in Theorem 3.14 is necessary if we just want to force any k-block, not necessarily one of size $\geq \delta + 1$.

If, in addition to being k-connected, G contains no triangle, the minimum degree needed to force a k-block comes down to $k+1$, and the k-blocks we find are also larger:

Theorem 3.15. *If G is k-connected, $\delta(G) \geq k+1$, and G contains no triangle, then either V is a k-block or G has at least two k-blocks. These can be chosen to be connected in G and of size at least $2\delta(G)$.*

Proof. Since $2\delta(G) > \frac{3}{2}k$, it suffices by Lemma 3.13 to show that $|A| \geq 2\delta(G)$ for every proper k-separation (A, B) of G. Pick a vertex $v \in A \setminus B$. As $d(v) \geq k+1$, it has a neighbour w in $A \setminus B$. Since v and w have no common neighbour, we deduce that $|A| \geq d(v) + d(w) \geq 2\delta(G)$. $\square$

Any k-connected, k-regular, triangle-free graph shows that the degree bound in Theorem 3.15 is sharp, because of the following observation:

Proposition 3.16. *If G is k-connected and k-regular, then G has no k-block unless $G = K^{k+1}$ (which contains a triangle).*

Proof. Suppose G has a k-block X. Pick a vertex $x \in X$. The k neighbours of x in G do not separate it from any other vertex of X, so all the other vertices of X are adjacent to x. But then X consists of precisely x and its k neighbours, since $|X| \geq k+1$. As this is true for every $x \in X$, it follows that $G = K^{k+1}$. $\square$

If we strengthen our regularity assumption to transitivity (i.e., assume that for every two vertices u, v there is an automorphism mapping u to v), then G has no k-blocks, regardless of its degree:

Theorem 3.17. *If $\kappa(G) = k \geq 1$ and G is vertex-transitive, then G has no k-block unless $G = K^{k+1}$.*

Proof. Unless G is complete (so that $G = K^{k+1}$), it has a proper k-separation. Hence V is not a k-block. Let us show that G has no k-block at all.

If G has a k-block, it has at least two, since V is not a k-block but every vertex lies in a k-block, by transitivity. Hence any tree-decomposition that distinguishes all the k-blocks of G has at least two parts. By Theorem 1.1, which we proved in Section 3.3, there exists such a tree-decomposition $(\mathcal{T}, \mathcal{V})$, which moreover has the property that every automorphism of G acts on the set of its parts. As $k \geq 1$, adjacent parts overlap in at least one vertex, so G has a vertex u that lies in at least two parts. But G also has a vertex v that lies in only one part (as long as no part of the decomposition contains another, which we may clearly assume): if t is a leaf of $\mathcal{T}$ and t' is its neighbour in $\mathcal{T}$, then every vertex in $V_t \setminus V_{t'}$ lies in no other part than V_t (see Section 3.1). Hence no automorphism of G maps u to v, a contradiction to the transitivity of G. $\square$

Theorems 3.14 and 3.17 together imply a well-known theorem of Mader [11] and Watkins [18], which says that every transitive graph of connectivity k has minimum degree at most $\frac{3}{2}k - 1$.

3.5 Forcing k-blocks by average degree conditions

As before, let us consider a non-empty graph $G = (V, E)$ fixed throughout this section. We denote its average degree by $d(G)$.

As remarked in the introduction, Mader [12] proved that if $d(G) \geq 4k$ then G has a $(k+1)$-connected subgraph. The vertex set of such a subgraph is k-inseparable, and hence extends to a k-block of G. Our first aim will be to show that if we seek to force a k-block in G directly, an average degree of $d(G) \geq 3k$ will be enough.

In the proof of that theorem, we may assume that G is a minimal with this property, so its proper subgraphs will all have average degrees smaller than $3k$. The following lemma enables us to utilize this fact. Given a set $S \subseteq V$, write $E(S, V)$ for the set of edges of G that are incident with a vertex in S.

Lemma 3.18. *If $\lambda > 0$ is such that $d(G) \geq 2\lambda > d(H)$ for every proper subgraph $H \neq \emptyset$ of G, then $|E(S, V)| > \lambda |S|$ for every set $\emptyset \neq S \subsetneq V$.*

Proof. Suppose there is a set $\emptyset \neq S \subsetneq V$ such that $|E(S, V)| \leq \lambda |S|$. Then our assumptions imply

$$|E(G - S)| = |E| - |E(S, V)| \geq \lambda |V| - \lambda |S| = \lambda |V \setminus S|,$$

so the proper subgraph $G - S$ of G contradicts our assumptions. $\square$

Theorem 3.19. *If $d(G) \geq 3k$, then G has a k-block. This can be chosen to be connected in G and of size at least $\delta(G) + 1 - k$.*

Proof. If $k = 0$, then the assertion follows directly. So we assume $k > 0$. Replacing G with a subgraph if necessary, we may assume that $d(G) \geq 3k$ but $d(H) < 3k$ for every proper subgraph H of G. By Lemma 3.18, this implies that $|E(S, V)| > \frac{3}{2}k|S|$ whenever $\emptyset \neq S \subsetneq V$; in particular, $\delta(G) > \frac{3}{2}k$.

Let $(\mathcal{T}, \mathcal{V})$ be a k-lean tree-decomposition of G, with $\mathcal{V} = (V_t)_{t\in\mathcal{T}}$ say. Pick a leaf t of $\mathcal{T}$. (If $\mathcal{T}$ has only one node, let t be this node.) Then $\mathring{V}_t \neq \emptyset$ by (T3), since V_t is not contained in any other part of $\mathcal{V}$.

If $|\mathring{V}_t| \leq k$ then, as also $|V_t \setminus \mathring{V}_t| \leq k$,

$$|E(\mathring{V}_t, V)| \leq \tfrac{1}{2}|\mathring{V}_t|^2 + k\,|\mathring{V}_t| \leq |\mathring{V}_t|\big(|\mathring{V}_t|/2 + k\big) \leq \tfrac{3}{2}k\,|\mathring{V}_t|,$$

which contradicts Lemma 3.18. So $|\mathring{V}_t| \geq k+1 \geq 2$. The set $\mathring{V}_t$ extends to a k-block $B \subseteq V_t$ with the desired properties as in the proof of Theorem 3.10. $\square$

Since our graph of Example 3.6 contains no k-block, its average degree is a strict lower bound for the minimum average degree that forces a k-block. By choosing the ladder in the construction of that graph long enough, we can make its average degree exceed $2k-1-\epsilon$ for any $\epsilon > 0$. The minimum average degree that will force a k-block thus lies somewhere between $2k-1$ and $3k$.

As we have seen, an average degree of $3k$ is sufficient to force a graph to contain a k-block. If we ask only that the graph should have a minor that contains a k-block, then a smaller average degree suffices:

Theorem 3.20. *If $d(G) \geq 2(k-1) > 0$, then G has a minor containing a k-block. This k-block can be chosen to be connected in the minor.*

Proof. Replacing G with a minor of itself if necessary, we may assume that $d(G) \geq 2(k-1)$ but $d(H) < 2(k-1)$ for every proper minor H of G. In particular, this holds for all subgraphs $\emptyset \neq H \subsetneq G$, so $\delta(G) \geq k$ by Lemma 3.18.

Let us show that any two adjacent vertices v and w have at least $k-1$ common neighbours. Otherwise, contracting the edge vw we lose one vertex and at most $k-1$ edges; as $|E|/|V| \geq k-1$ by assumption, this ratio (and hence the average degree) will not decrease, contradicting the minimality of G.

Let $(\mathcal{T}, \mathcal{V})$ be a k-lean tree-decomposition of G, with $\mathcal{V} = (V_t)_{t\in\mathcal{T}}$ say, and let t be a leaf of $\mathcal{T}$. (If $\mathcal{T}$ has only one node, let t be this node.) We shall prove that V_t is k-inseparable, and hence a k-block, in G.

As $(\mathcal{T}, \mathcal{V})$ is k-lean, every vertex $a \in A_t := V_t \cap \bigcup_{t'\neq t} V_{t'}$ has a neighbour v in $\mathring{V}_t$, as otherwise $X := A_t \setminus \{a\}$ would separate A_t from every set $X \cup \{v\}$ with $v \in \mathring{V}_t$, which contradicts k-leanness since $|X \cup \{v\}| = |A_t| \leq k$. As a and v have $k-1$ common neighbours in G, which must lie in V_t, we find that every vertex in A_t, and hence every vertex of V_t, has at least k neighbours in V_t.

As $\mathring{V}_t \neq \emptyset$ and hence $|V_t| \geq \delta(G) + 1 \geq k+1$, it suffices to show that two vertices $u, v \in V_t$ can never be separated in G by $\leq k$ other vertices. But this follows from k-leanness: pick a set N_u of k neighbours of u in V_t and a set N_v of k neighbours of v in V_t to obtain two $(k+1)$-sets $N_u \cup \{u\}$ and $N_v \cup \{v\}$ that are joined in G by $k+1$ disjoint paths; hence u and v cannot be separated by $\leq k$ vertices. $\square$

Recall that the graphs of Example 3.6 have average degrees of at least $2k-1-\epsilon$. So these graphs show that obtaining a k-block in G is indeed harder than obtaining a k-block in a minor of G, which these graphs must have by Theorem 3.20. (And they do: they even have $K^{3k/2}$-minors.)

Chapter 4

Profiles and connectivity systems

In this chapter we establish the main methods, which we later use in Chapter 5 to prove most of the main results stated in the Introduction. In Section 4.1 we introduce the notion of a profile and show how those profiles relate to nested separation systems and the tree-decomposition obtained from those separation systems. In Section 4.2 we show how to distinguish certain sets of profiles, in a simple way, by means of a nested separation system. In Section 4.3 we refine these methods and show that there are several different strategies to distinguish a set of profiles. An essential part of Section 4.3 is based on an early version of [2]. In Section 4.4 we adapt the notion of a connectivity system, as defined by Geelen et al. in [9], to our notion of a separation from Chapter 2. We then show that the k-profiles of such a connectivity system can be pairwise distinguished by means of a tree-decomposition. We conclude the section and the chapter by presenting further results of [2], which establish some bounds on the needed number of parts in a tree-decomposition distinguishing all the k-profiles. We have presented the main ideas of this chapter, except for those found in [2], in a more abstract setup in [10].

4.1 Profiles

One of the main results of Chapter 3 is that all the k-blocks of a graph G can be distinguished by a single tree-decomposition $(\mathcal{T}, \mathcal{V})$ of G of adhesion at most k. That is, distinct k-blocks are contained in different parts of $(\mathcal{T}, \mathcal{V})$. By Theorem 2.17 (ii) those parts of $(\mathcal{T}, \mathcal{V})$ that host a k-block must be large $\mathcal{N}$-blocks, where $\mathcal{N}$ is the (nested) set of separations of G induced by $(\mathcal{T}, \mathcal{V})$. In fact, it was essential to our proof that we were able to identify both the k-blocks of G and those parts of $(\mathcal{T}, \mathcal{V})$ that host them, as large $\mathcal{S}$-blocks for an appropriate system $\mathcal{S}$ of separations of G.

However, to extend our approach to tangles, of both graphs and matroids, the concept of an $\mathcal{S}$-block is no longer sufficient. This has two reasons. First, even in a graph G, we need not be able to identify all tangles of G as $\mathcal{S}$-blocks for some separation system $\mathcal{S}$ of G. And second, in a tree-decomposition of a matroid M there may be 'essential' hub-nodes: a part of $(\mathcal{T}, \mathcal{V})$ may be empty and still host a tangle of large order (cp. Section 5.3). Obviously, such a part is not an $\mathcal{S}$-block for any separation system $\mathcal{S}$ of M.

We solve this problem by introducing a new notion of $\mathcal{S}$-inseparable 'object', for $\mathcal{S}$ being a system of separations of some finite set V, which we describe, in analogy to the definition of a tangle by Robertson and Seymour [14], as an 'orientation' of $\mathcal{S}$ instead of a subset of V. We will develop this new notion guided by the notion of a large $\mathcal{S}$-block.

For a system $\mathcal{S}$ of separations of a finite set V and a large $\mathcal{S}$-block $b \subseteq V$ consider the following set $P_{\mathcal{S}}(b)$ of separations:

$$P_{\mathcal{S}}(b) := \{(A,B) \in \mathcal{S} \mid b \subseteq B\}. \tag{4.1}$$

By definition of a large $\mathcal{S}$-block, $P_{\mathcal{S}}(b)$ contains precisely one of $\{(A,B),(B,A)\}$ for all $(A,B) \in \mathcal{S}$. In general, we call a set of separations which contains precisely one of $\{(A,B),(B,A)\}$ for all $(A,B) \in \mathcal{S}$, and no other separations, an *orientation* of $\mathcal{S}$. In contrast to most arbitrary orientations of $\mathcal{S}$, the choices that $P = P_{\mathcal{S}}(b)$ makes on $\mathcal{S}$ are consistent in the following sense:

(P1) For every $(A,B) \in P$ and every separation $(C,D) \leq (A,B)$ we have $(D,C) \notin P$;

(P2) For all $(A,B),(C,D) \in P$ we have $(B \cap D, A \cup C) \notin P$.

Let us call an arbitrary set of separations a *profile* if it satisfies (P1) and (P2). Given a set $\mathcal{S}$ of separations, an orientation of $\mathcal{S}$ that also is a profile is called an *$\mathcal{S}$-profile.* So for a large $\mathcal{S}$-block b the set $P_{\mathcal{S}}(b)$ forms an $\mathcal{S}$-profile, which we call the *$\mathcal{S}$-profile of b.* Given an $\mathcal{S}$-profile P we call it an *$\mathcal{S}$-block profile* if there exists a large $\mathcal{S}$-block b such that $P_{\mathcal{S}}(b) = P$. If $b \neq b'$ are different large $\mathcal{S}$-blocks, then by Lemma 2.1 there is a separation $(A,B) \in \mathcal{S}$ such that $b \subseteq A$ and $b' \subseteq B$. Hence we have $(A,B) \in P_{\mathcal{S}}(b) \setminus P_{\mathcal{S}}(b')$ and $(B,A) \in P_{\mathcal{S}}(b') \setminus P_{\mathcal{S}}(b)$, such that $P_{\mathcal{S}}(b)$ and $P_{\mathcal{S}}(b')$ are distinct (indeed incomparable under set inclusion). That is, an $\mathcal{S}$-block profile P satisfies $P = P_{\mathcal{S}}(b)$ for precisely one large $\mathcal{S}$-block b, which we then denote by $b(P)$, and we say that P and $b(P)$ *correspond.*

Now we are able to uniquely identify a large $\mathcal{S}$-block by means of a an $\mathcal{S}$-profile. However, not every $\mathcal{S}$-profile arises in that way from an $\mathcal{S}$-block. We shall see later that 'most' hub-nodes of a tree-decomposition obtained from a nested separation system $\mathcal{N}$ will define an $\mathcal{N}$-profile, which then clearly is distinct to all $\mathcal{N}$-block profiles (since those correspond to 'large' block-nodes). However, not all hub-nodes give rise to an $\mathcal{N}$-profile. In order to describe all parts of a tree-decomposition in terms of (the corresponding separation system) $\mathcal{N}$, we need an even weaker notion than that of a profile. Let us call a set of separations that satisfies (P1) – but not necessarily (P2) – a *preference.* And, as for profiles, we define an *$\mathcal{S}$-preference* to be a preference that is an $\mathcal{S}$-orientation.

Now, for a nested separation system $\mathcal{N}$, this notion of an $\mathcal{N}$-preference corresponds precisely to the parts of the tree-decomposition obtained from $\mathcal{N}$, as follows. As described in Chapter 2, each orientation of an edge $e = \{t_1,t_2\}$, that is one of the two ordered pairs (t_1,t_2) or (t_2,t_1), of the decomposition tree $\mathcal{T}$ of a tree-decomposition $(\mathcal{T},\mathcal{V})$ of V induces a separation of V. Furthermore, it is obvious that every node $t \in \mathcal{T}$ can be described as a unique orientation of the edges of $\mathcal{T}$: for every edge $e = \{t_1,t_2\}$ choose (t_{3-i},t_i) such that t_i and t are in the same component of $\mathcal{T} - e$ (we make e 'point towards t'). This orientation then gives rise to a set of separations of V, which we shall denote by $P_{\mathcal{N}}(t)$. We then say that the set $P_{\mathcal{N}}(t)$ of separations and the part V_t *correspond*. Now we can state and proof the before-mentioned correspondence between the $\mathcal{N}$-preferences and the parts of $(\mathcal{T}, \mathcal{V})$ in a formal way:

Proposition 4.1. *Let $\mathcal{N}$ be the nested separation system induced by a tree-decomposition $(\mathcal{T}, \mathcal{V})$ of a finite set V.*

(i) *Every $\mathcal{N}$-preference corresponds to a unique part of $(\mathcal{T}, \mathcal{V})$;*

(ii) *Every part of $(\mathcal{T}, \mathcal{V})$ corresponds to a unique $\mathcal{N}$-preference.*

Proof. A subset $P \subseteq \mathcal{N}$ which contains precisely one of $\{(A, B), (B, A)\}$ for every $(A, B) \in \mathcal{N}$ defines an orientation $\mathcal{O}_P$ of the edges of $\mathcal{T}$. We say that a node $t \in \mathcal{N}$ has *out-degree k in $\mathcal{O}_P$* if there are precisely k other nodes s of $\mathcal{T}$ with $(t, s) \in \mathcal{O}_P$. A node $t \in \mathcal{T}$ is a *sink of $\mathcal{O}_P$* if it has out-degree 0 in $\mathcal{O}_P$. Then P corresponds to the part V_t of $(\mathcal{T}, \mathcal{V})$, if and only if t is the unique sink of $\mathcal{O}_P$. It is easy to see that $\mathcal{O}_P$ has a unique sink, if and only if every node of $\mathcal{T}$ has out-degree at most 1 in $\mathcal{O}_P$. This, however, is the case, if and only if P satisfies (P1), which implies both (i) and (ii). $\square$

Once more, consider a large $\mathcal{S}$-block b for some separation system $\mathcal{S}$. Then b is clearly contained in a large $\mathcal{N}$-block for every nested separation system $\mathcal{N} \subseteq \mathcal{S}$, in particular b will be contained in, or *inhabit*, a part of the tree-decomposition $(\mathcal{T}, \mathcal{V})$ obtained from $\mathcal{N}$. Now we would like to express the fact that b inhabits a part of $(\mathcal{T}, \mathcal{V})$ in terms of the $\mathcal{N}$-profile of b. In that way we are able to extend our terminology to arbitrary profiles.

Let $\mathcal{N}$ be a nested separation system and let $(\mathcal{T}, \mathcal{V})$ be the tree-decomposition obtained from $\mathcal{N}$. A profile P *inhabits* a part V_t of $(\mathcal{T}, \mathcal{V})$ if its intersection with $\mathcal{N}$ extends to a unique $\mathcal{N}$-preference, which corresponds to V_t. That is, by Proposition 4.1, P inhabits V_t if we have $P \cap \mathcal{N} \subseteq P_{\mathcal{N}}(t)$ and $P \cap \mathcal{N} \not\subseteq P_{\mathcal{N}}(t')$, for every $t' \neq t$. By Lemma 4.3 we obtain a characterisation for when this is the case. Let us first emphasise a trivial but fundamental fact about profiles:

$$\textit{Every subset of a profile is a profile.} \tag{4.2}$$

One direction of Lemma 4.3 will be a special case of the following more general lemma, which shows that an $\mathcal{S}$-profile P is uniquely determined by its maximal elements – those separations that are maximal in P with respect to $\leq$. In its proof we make use of the following fact, which follows easily from the definition of an $\mathcal{S}$-profile.

$$\textit{Distinct } \mathcal{S}\textit{-profiles are incomparable under set inclusion.} \tag{4.3}$$

Now we are able to proof the before-mentioned lemma.

Lemma 4.2. *Let $\mathcal{S}$ be an arbitrary separation system and let P be an $\mathcal{S}$-profile. If $P' \subseteq \mathcal{S}$ is an arbitrary profile that contains all maximal elements of P, then P is the unique $\mathcal{S}$-profile to which P' extends.*

Proof. Assume first that P' contains all maximal elements of P. Suppose that there is a separation $(C, D) \in P' \setminus P$. Since P is an $\mathcal{S}$-profile we then have

$(D,C) \in P$. By assumption (C,D) is not maximal in P, such that there is $(A,B) \in P$ with $(C,D) \leq (A,B)$, which yields $(D,C) \notin P$ by (P1), a contradiction. Therefore we have $P' \subseteq P$. Now suppose there is an $\mathcal{S}$-profile $Q \neq P$ and suppose we have $P' \subseteq Q$. By the previous argument (for $P' = Q$) we then have $Q \subseteq P$, in contradiction to (4.3). □

The converse of Lemma 4.2 need not hold in this generality. However, if we consider a nested separation system $\mathcal{N}$ then it does. As before let $(\mathcal{T}, \mathcal{V})$ be the tree-decomposition obtained from $\mathcal{N}$.

Lemma 4.3. *A profile P inhabits a part V_t if and only if P contains all maximal separations in $P_{\mathcal{N}}(t)$.*

Proof. The 'if'-direction is due to Lemma 4.2 (applied to $P \cap \mathcal{N}$). For the other direction consider a profile P that inhabits a part V_t. By definition this means that we have $P \cap \mathcal{N} \subseteq P_{\mathcal{N}}(t)$ and $P \cap \mathcal{N} \not\subseteq P_{\mathcal{N}}(t')$, for every $t' \neq t$. Suppose that there is a maximal element (A,B) in $P_{\mathcal{N}}(t)$ that is not contained in P. Then it is straightforward to check that $P_{\mathcal{N}}(t) \smallsetminus \{(A,B)\} \cup \{(B,A)\}$ is an $\mathcal{N}$-preference, which by Proposition 4.1 corresponds to a part $V_{t'}$ with $t \neq t'$, contradicting our assumption. □

Given a separation system $\mathcal{S}$, let us say that a profile P *orients* $\mathcal{S}$ if $P \cap \mathcal{S}$ is an orientation of $\mathcal{S}$. We obtain the following corollary:

Corollary 4.4. *Every profile that orients a nested separation system $\mathcal{N}$ inhabits a unique part of the tree-decomposition obtained from $\mathcal{N}$.* □

4.2 Distinguishing profiles

Given two profiles P and P' we say that a separation (A,B) *distinguishes* P and P' if we have $(A,B) \in P$ and $(B,A) \in P'$, or vice versa. A set $\mathcal{S}$ of separations *distinguishes* P and P' if it contains a separation that does so. The set $\mathcal{S}$ is said to *distinguish* a set $\mathcal{P}$ of profiles if it distinguishes each pair of profiles in $\mathcal{P}$, and $\mathcal{P}$ is then called *$\mathcal{S}$-distinguishable*.

Our aim in this section is to find inside a given separation system $\mathcal{S}$ a nested subsystem $\mathcal{N}$ which still distinguishes a given set $\mathcal{P}$ of $\mathcal{S}$-distinguishable profiles that all orient $\mathcal{S}$. As we have seen in Section 2.4 this need not succeed for all choices of $\mathcal{S}$ and $\mathcal{P}$ – not even in the special case, when $\mathcal{P}$ is a set of $\mathcal{S}$-block profiles.

However, as in Section 2.4 we can give a sufficient condition on $\mathcal{S}$ and $\mathcal{P}$ under which we will always find such a nested subsystem. We say that a profile P has a *tendency* towards the corner $B \cap D$ of $\{(A,B),(C,D)\}$ if we have $(A,B),(C,D) \in P$. Now (P2) says that a profile will never contain the inverse of a corner-separation corresponding to a corner towards which it has a tendency, but it need not contain the corner-separation itself. The following condition requieres that for each pair P, P' of profiles in $\mathcal{P}$ which have a tendency towards opposite corners of a pair of crossing separations in $\mathcal{S}$ there must exist a separation in $P \cap \mathcal{S}$ (and by symmetry also one in $P' \cap \mathcal{S}$) that is larger than the corresponding corner-separation. Given a system $\mathcal{S}$ of separations and a set $\mathcal{P}$

of profiles that all orient $\mathcal{S}$ we say that $\mathcal{S}$ *separates* $\mathcal{P}$ *well* if the following holds for every pair of crossing separations $(A,B),(C,D) \in \mathcal{S}$:

> *For all $P, P' \in \mathcal{P}$ with $(A,B),(C,D) \in P$ and $(B,A),(D,C) \in P'$ there is an $(E,F) \in P \cap \mathcal{S}$ with $(A \cup C, B \cap D) \leq (E,F)$.*

Note that this separation (E,F) will also distinguish P from P': by assumption we have $(E,F) \in P$ and by (2.2) we have $(F,E) \leq (B \cap D, A \cup C) \leq (B,A)$ such that (P1) and the fact that P' orients $\mathcal{S}$ yield $(F,E) \in P'$. Let us call a separation $\mathcal{P}$-*relevant* if it distinguishes some pair of profiles in $\mathcal{P}$. Then all the separations (A,B), (C,D), (E,F), mentioned above, are $\mathcal{P}$-relevant. Let $\mathcal{R}$ denote the set of all $\mathcal{P}$-relevant separations in $\mathcal{S}$. Then we obviously have

$$\mathcal{S} \textit{ separates } \mathcal{P} \textit{ well if and only if } \mathcal{R} \textit{ separates } \mathcal{P} \textit{ well.} \tag{4.4}$$

That is, in many situations we may directly assume $\mathcal{S}$ to be $\mathcal{P}$-relevant. Now consider a separation system $\mathcal{S}$ and a non-empty set $\mathcal{P}$ of profiles that all orient $\mathcal{S}$. The pair $(\mathcal{S},\mathcal{P})$ is a *task* if $\mathcal{P}$ is $\mathcal{S}$-distinguishable and $\mathcal{S}$ separates $\mathcal{P}$ well. If $(\mathcal{S},\mathcal{P})$ and $(\mathcal{S}',\mathcal{P}')$ are tasks with $\mathcal{S}' \subseteq \mathcal{S}$ and $\mathcal{P}' \subseteq \mathcal{P}$ then the latter is called a *subtask* of the first. A task is called *reduced* if every separation in $\mathcal{S}$ is $\mathcal{P}$-relevant. Given a task $(\mathcal{S},\mathcal{P})$ let $\mathcal{R}$, as above, denote the set of all $\mathcal{P}$-relevant separations in $\mathcal{S}$, then we call $(\mathcal{R},\mathcal{P})$ the *reduction* of $(\mathcal{S},\mathcal{P})$, which is a reduced task due to (4.4).

Recall that a separation (E,F) is called extremal in $\mathcal{S}$ if for all $(C,D) \in \mathcal{S}$ we have $(C,D) \leq (E,F)$ or $(D,C) \leq (E,F)$. So in particular, every extremal separation is nested with every other separation in $\mathcal{S}$. A separation is called *extremal* in a task $(\mathcal{S},\mathcal{P})$, if it is extremal in $\mathcal{S}$. The following lemma will be the cornerstone in the proof of our main result of this section.

Lemma 4.5. *If $(\mathcal{S},\mathcal{P})$ is a reduced task then every $\leq$-maximal separation in $\mathcal{S}$ is extremal in $(\mathcal{S},\mathcal{P})$.*

Proof. Let (A,B) be a maximal separation in $\mathcal{S}$ and consider an arbitrary separation $(C,D) \in \mathcal{S}$. If (A,B) and (C,D) are nested, then we have $(C,D) \leq (A,B)$ or $(D,C) \leq (A,B)$ by the maximality of (A,B). So suppose we have $(A,B) \nparallel (C,D)$. Since $\mathcal{S}$ is $\mathcal{P}$-relevant there are $P, P' \in \mathcal{P}$ which have a tendency towards opposite corners of $\{(A,B),(C,D)\}$. Renaming (C,D) as (D,C) if necessary we may assume $(A,B),(C,D) \in P$ and $(B,A),(D,C) \in P'$. Since $\mathcal{S}$ separates $\mathcal{P}$ well this yields a separation $(E,F) \in P \cap \mathcal{S}$ with $(A \cup C, B \cap D) \leq (E,F)$. By (2.4) this yields $(A,B) \leq (E,F)$ and $(C,D) \leq (E,F)$. So the maximality of (A,B) gives $(C,D) \leq (E,F) = (A,B)$, a contradiction. □

We further have:

Lemma 4.6. *If $(\mathcal{S},\mathcal{P})$ is a reduced task and $(A,B) \in \mathcal{S}$ extremal in $(\mathcal{S},\mathcal{P})$, then there is a unique profile $P_{(A,B)}$ in $\mathcal{P}$ with $(A,B) \in P_{(A,B)}$.*

Proof. Let (A,B) be extremal in $(\mathcal{S},\mathcal{P})$. Since $\mathcal{S}$ is $\mathcal{P}$-relevant there is a profile P in $\mathcal{P}$ with $(A,B) \in P$. Suppose that there is another profile $P' \neq P$ in $\mathcal{P}$ which also contains (A,B). Since $\mathcal{P}$ is $\mathcal{S}$-distinguishable there is $(C,D) \in \mathcal{S}$ with $(C,D) \in P$ and $(D,C) \in P'$. Since (A,B) is extremal in $\mathcal{S}$ we have either $(C,D) \leq (A,B)$ or $(D,C) \leq (A,B)$, we may assume the former. But then (P1) yields $(D,C) \notin P'$, a contradiction. Hence, $P_{(A,B)} := P$ is the unique profile in $\mathcal{P}$ containing (A,B). □

We are now ready to prove the first main result of this section. The proof follows essentially the same approach as the proof of Theorem 2.19. The interested reader may compare the two proofs in order to see how to 'translate' from $\mathcal{S}$-inseparable sets to profiles that orient $\mathcal{S}$.

Theorem 4.7. *If $(\mathcal{S}, \mathcal{P})$ is a task, then $\mathcal{S}$ has a $\mathcal{P}$-relevant nested subsystem $\mathcal{N}(\mathcal{S}, \mathcal{P})$ that distinguishes each pair of profiles in $\mathcal{P}$.*

Proof. The proof is by induction on $p := |\mathcal{P}|$. For $p \leq 1$ we do not have to distinguish any profiles, so we set $\mathcal{N}(\mathcal{S}, \mathcal{P}) = \emptyset$, which clearly has all stated properties.

For $p \geq 2$ let $\mathcal{R} \subseteq \mathcal{S}$ be the set of all $\mathcal{P}$-relevant separations in $\mathcal{S}$. Since $\mathcal{P}$ is $\mathcal{S}$-distinguishable and contains at least two distinct profiles we know that $\mathcal{R}$ is not empty. Hence, Lemma 4.5 yields that there are extremal elements in $\mathcal{R}$. Let $\mathcal{E} \subseteq \mathcal{R}$ be the set of all those separations that are extremal in $\mathcal{R}$, put

$$\mathcal{P}_{\mathcal{E}} := \{P_{(A,B)} \mid (A, B) \in \mathcal{E}\},$$

and let $\mathcal{P}' := \mathcal{P} \setminus \mathcal{P}_{\mathcal{E}}$. Since $\mathcal{P}'$ is a subset of $\mathcal{P}$, we know that $\mathcal{S}$ separates $\mathcal{P}'$ well. If $\mathcal{P}'$ is non-empty, then $(\mathcal{S}, \mathcal{P}')$ is a task and by the induction hypothesis we obtain a nested $\mathcal{P}'$-relevant subsystem $\mathcal{N}(\mathcal{S}, \mathcal{P}') \subseteq \mathcal{R}$ which separates each pair of profiles in $\mathcal{P}'$. If $\mathcal{P}'$ is empty, put $\mathcal{N}(\mathcal{S}, \mathcal{P}') := \emptyset$. Now let

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

By definition of extremality and symmetry of nestedness all separations in $\overline{\mathcal{E}}$ are nested with all separations in $\mathcal{R}$. Therefore

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

is a nested subsystem of $\mathcal{S}$, in fact it is a subsystem of $\mathcal{R}$, such that $\mathcal{N}(\mathcal{S}, \mathcal{P})$ is $\mathcal{P}$-relevant. Let us check that $\mathcal{N}(\mathcal{S}, \mathcal{P})$ distinguishes each pair of profiles in $\mathcal{P}$. Consider distinct profiles $P, P' \in \mathcal{P}$. If both lie in $\mathcal{P}'$ then they are distinguished by $\mathcal{N}(\mathcal{S}, \mathcal{P}')$, which is a subsystem of $\mathcal{N}(\mathcal{S}, \mathcal{P})$. If one of them, say P, lies in $\mathcal{P}_{\mathcal{E}}$, then we have $P = P_{(A,B)}$ for some $(A, B) \in \mathcal{E}$. By Lemma 4.6 we then have $(B, A) \in P'$, such that (A, B) distinguishes P and P'. Since (A, B) lies in $\mathcal{N}(\mathcal{S}, \mathcal{P})$ this finishes the proof. $\square$

With the same argument as in Section 2.4 after Theorem 2.19 (cp. Page 23), the nested separation system we construct in the proof of Theorem 4.7 is *canonical* in the following sense.

Corollary 4.8. *Let $\mathcal{S}$ and $\mathcal{P}$ be as in Theorem 4.7, and let $\mathcal{N}(\mathcal{S}, \mathcal{P})$ be the nested subsystem of $\mathcal{S}$ constructed in the proof. Then for every permutation α of V we have $\mathcal{N}(\mathcal{S}^\alpha, \mathcal{P}^\alpha) = \mathcal{N}(\mathcal{S}, \mathcal{P})^\alpha$. In particular, if $\mathcal{S}$ and $\mathcal{P}$ are invariant under the action of a group Γ of permutations of V, then so is $\mathcal{N}(\mathcal{S}, \mathcal{P})$.* $\square$

4.3 Different strategies to distinguish profiles

An essential part of this section is based on an early version of [2]. While 'being canonical' is a strong property of the nested separation systems $\mathcal{N}(\mathcal{S}, \mathcal{P})$ that we have constructed in the previous section, it does not uniquely determine those systems in terms of $\mathcal{S}$ and $\mathcal{P}$. In this section we show that we can modify our approach in different ways, each of which produces a canonical separation system as stated in Theorem 4.7. In the end of this section we shall see that these different approaches may in fact yield distinct canonical separation systems for given $\mathcal{S}$ and $\mathcal{P}$.

A profile P is called *extremal* in $(\mathcal{S}, \mathcal{P})$ if there is an extremal separation $(A, B) \in \mathcal{S}$ with $P = P_{(A,B)}$ (as given by Lemma 4.6). It is not too complicated to show directly that every extremal profile P is well-separated, where a profile P is called *well-separated* in $\mathcal{S}$ if it has the following property:

$$\text{For every } (A,B) \nparallel (C,D) \in \mathcal{S} \text{ such that } (A,B),(C,D) \in P \text{ there exists } (E,F) \in P \cap \mathcal{S} \text{ with } (A \cup C, B \cap D) \le (E,F). \tag{4.5}$$

However, instead of proving that extremal profiles are well-separated in $\mathcal{S}$, we give by the following lemma a characterization of well-separated profiles, from which this is easily obtained.

Lemma 4.9. *A profile P that orients a separation system $\mathcal{S}$ is well-separated in $\mathcal{S}$ if and only if the maximal elements of $(P \cap \mathcal{S}, \le)$ are nested with every separation in $\mathcal{S}$.*

Proof. Consider a profile P such that every $\le$-maximal element of $(P \cap \mathcal{S})$ is nested with every separation in $\mathcal{S}$ – we refer to this property by $(*)$. Now consider two separations $(A, B) \nparallel (C, D) \in \mathcal{S}$ such that we have $(A, B), (C, D) \in P$. By $(*)$ neither (A, B) nor (C, D) can be maximal in $(P \cap \mathcal{S}, \le)$. So there is a maximal element (E, F) in $(P \cap \mathcal{S}, \le)$ with $(A, B) \le (E, F)$. Again by $(*)$ we have $(E, F) \parallel (C, D)$, which by definition means that (E, F) is $\le$-comparable with either (C, D) or (D, C). Since (A, B) is not nested with (C, D) we have $(E, F) \not\le (C, D)$ and $(E, F) \not\le (D, C)$, and since both (C, D) and (E, F) are in P, axiom (P1) yields $(D, C) \not\le (E, F)$. Hence, we have $(C, D) \le (E, F)$ and therefore we have $(A \cup C, B \cap D) \le (E, F)$, which proves P to be well-separated.

To show the other direction, let P be a well-separated profile and $(A, B) \in P$ maximal in $(P \cap \mathcal{S}, \le)$. Suppose there is $(C, D) \in \mathcal{S}$ with $(A, B) \nparallel (C, D)$. Since P orients $\mathcal{S}$, and by symmetry of nestedness, we may assume $(C, D) \in P$. Then by (4.5) there is $(E, F) \in P \cap \mathcal{S}$ with $(A \cup C, B \cap D) \le (E, F)$, and we have $(A, B) \le (E, F)$ as well as $(C, D) \le (E, F)$. But then $(E, F) = (A, B)$ by maximality of (A, B) and therefore $(A, B) \parallel (C, D)$, in contradiction to our assumption. So P satisfies $(*)$, which finishes the proof. $\square$

We say that a separation (A, B) is *locally extremal* in $(\mathcal{S}, \mathcal{P})$ if there is a well-separated profile $P \in \mathcal{P}$ such that (A, B) is maximal in $(P \cap \mathcal{S}, \le)$. By Lemma 4.9 we hence know that every separation that is locally extremal in $(\mathcal{S}, \mathcal{P})$ is nested with every separation in $\mathcal{S}$. Separations with this property seem to be a good start for *solving* our task – to find a nested subsystem $\mathcal{N} \subseteq \mathcal{S}$ distinguishing all the profiles in $\mathcal{P}$ – since they will not prevent us from choosing any other separation of $\mathcal{S}$ later on. Let us denote the symmetric closure of the set of all separations in $\mathcal{S}$ that are: extremal in $(\mathcal{S}, \mathcal{P})$ by $\mathsf{ext}(\mathcal{S}, \mathcal{P})$, locally extremal in $(\mathcal{S}, \mathcal{P})$ by $\mathsf{loc}(\mathcal{S}, \mathcal{P})$, nested with every other separation in $\mathcal{S}$ by $\mathsf{max}(\mathcal{S}, \mathcal{P})$. We have the following relation between those nested separation systems:

$$\mathsf{ext}(\mathcal{S}, \mathcal{P}) \subseteq \mathsf{max}(\mathcal{S}, \mathcal{P}) \supseteq \mathsf{loc}(\mathcal{S}, \mathcal{P}). \tag{4.6}$$

In general, there need not be any relation between the extremal and the locally extremal separations of a task. In a reduced task, however, every extremal separation is also locally extremal (witnessed by the corresponding extremal profile), while the converse is not necessarily true. As we have seen earlier we can always obtain a reduced subtask $(\mathcal{R}, \mathcal{P})$ of $(\mathcal{S}, \mathcal{P})$ with the same set of profiles. For solving a task it may obviously make sense to consider $\mathsf{ext}(\mathcal{R}, \mathcal{P})$ instead of $\mathsf{ext}(\mathcal{S}, \mathcal{P})$: in order to distinguish some profiles in $\mathcal{P}$ it is sufficient to consider only separations in $\mathcal{R}$, and if solving our task is not trivial (if $|\mathcal{P}| \geq 2$) then $\mathsf{ext}(\mathcal{R}, \mathcal{P})$ will not be empty, by Lemma 4.5. To ease up notation later let us define the following abbreviations:

- $\mathsf{ext}_{\mathcal{R}}(\mathcal{S}, \mathcal{P}) := \mathsf{ext}(\mathcal{R}, \mathcal{P})$;
- $\mathsf{loc}_{\mathcal{R}}(\mathcal{S}, \mathcal{P}) := \mathsf{loc}(\mathcal{R}, \mathcal{P})$;
- $\mathsf{max}_{\mathcal{R}}(\mathcal{S}, \mathcal{P}) := \mathsf{max}(\mathcal{R}, \mathcal{P})$.

We then have

$$\emptyset \neq \mathsf{ext}_{\mathcal{R}}(\mathcal{S}, \mathcal{P}) \subseteq \mathsf{loc}_{\mathcal{R}}(\mathcal{S}, \mathcal{P}) \subseteq \mathsf{max}_{\mathcal{R}}(\mathcal{S}, \mathcal{P}). \tag{4.7}$$

In the following we consider each of $\{\mathsf{ext}, \mathsf{loc}, \mathsf{max}, \mathsf{ext}_{\mathcal{R}}, \mathsf{loc}_{\mathcal{R}}, \mathsf{max}_{\mathcal{R}}\}$ as a function that maps a task $(\mathcal{S}, \mathcal{P})$ to a nested subsystem of $\mathcal{S}$. The way we will proceed is as follows: starting with a task $(\mathcal{S}, \mathcal{P})$ we shall apply one of the previously defined functions to obtain a nested subsystem $\mathcal{N}$ of $\mathcal{S}$, and then we aim to 'split' those $\mathcal{N}$-profiles that extend to at least two distinct profiles in $\mathcal{P}$.

Given a nested separation system $\mathcal{N}$ and a separation (A, B) that is nested with $\mathcal{N}$, we say that (A, B) *splits* an $\mathcal{N}$-preference P, if both $P \cup \{(A, B)\}$ and $P \cup \{(B, A)\}$ are preferences.

Lemma 4.10. *Let $\mathcal{N}$ be a nested separation system. Every proper separation $(A, B) \notin \mathcal{N}$ that is nested with $\mathcal{N}$ splits a unique $\mathcal{N}$-preference.*

Proof. Since (A, B) is nested with $\mathcal{N}$, we have for every separation $(C, D) \in \mathcal{N}$ that either (C, D) or (D, C) is less than either (A, B) or (B, A). We show that

$$P := \{(C, D) \in \mathcal{N} \mid (C, D) \leq (A, B)\} \cup \{(E, F) \in \mathcal{N} \mid (E, F) \leq (B, A)\}$$

is a profile. From (2.2) and (2.3) we obtain that P is an orientation of $\mathcal{N}$.

To check (P1) consider a separation $(C, D) \in P$ and $(E, F) \in \mathcal{N}$ such that $(E, F) \leq (C, D)$. By definition of P we have either $(C, D) \leq (A, B)$ or $(C, D) \leq (B, A)$, let us assume the former. Then by transitivity we have $(E, F) \leq (A, B)$ which implies $(E, F) \in P$, and therefore $(F, E) \notin P$.

Suppose for a contradiction that there is another $\mathcal{N}$-preference $P' \neq P$ such that (A, B) splits P'. Then there is $(C, D) \in \mathcal{N}$ such that $(C, D) \in P$ and $(D, C) \in P'$. By symmetry and the construction of P we may assume $(C, D) \leq (A, B)$. But then $P' \cup \{(A, B)\}$ is not a preference, since (D, C) is contained in P'. This contradiction finishes the proof. $\square$

This is the right place to prove the following lemma, though we will not make use of it before the proof of Theorem 4.20 in Section 4.4.

Lemma 4.11. *Let $\mathcal{N}$ be a nested separation system and let P be an $\mathcal{N}$-preference. Let $\mathcal{N}' \supseteq \mathcal{N}$ be a nested separation system such that no separation in $\mathcal{N}' \setminus \mathcal{N}$ splits P, then P extends to a unique $\mathcal{N}'$-preference.*

Proof. Let $(C, D) \in \mathcal{N}' \smallsetminus \mathcal{N}$. By assumption (C, D) does not split P, so we may assume that $P \cup \{(C, D)\}$ is not a preference. This implies that there is $(A, B) \in P$ such that $(D, C) \leq (A, B)$, and therefore $(B, A) \leq (C, D)$. But then $P \cup \{(D, C)\}$ is a preference. Suppose not. Then there is $(E, F) \in P$ such that $(C, D) \leq (E, F)$, which implies $(B, A) \leq (E, F)$. This is in contradiction to (P1) since we have $(A, B) \in P$.

Let $\mathcal{M} := \mathcal{N} \cup \{(C, D), (D, C)\}$. Then $P' := P \cup \{(D, C)\}$ is the unique $\mathcal{M}$-preference to which P extends and the assertion of the lemma follows by induction on $|\mathcal{N}' \smallsetminus \mathcal{N}|$. $\square$

For the ongoing of this section we need another lemma.

Lemma 4.12. *Let $\mathcal{N}$ be a nested separation system. If two separations (A, B) and (C, D) split distinct $\mathcal{N}$-preferences, then (A, B) and (C, D) are nested.*

Proof. By Lemma 4.10 and the assumptions in the statement, there are unique $\mathcal{N}$-preferences $P \neq P'$ such that (A, B) splits P and (C, D) splits P'. Let $(E, F) \in \mathcal{N}$ such that $(E, F) \in P$ and $(F, E) \in P'$. By symmetry of nestedness and by what we know about P and P' from the proof of Lemma 4.10, we may assume $(E, F) \leq (A, B)$ and $(F, E) \leq (C, D)$. Then due to (2.2) we have $(B, A) \leq (F, E) \leq (C, D)$, which implies that (A, B) and (C, D) are nested with each other. $\square$

Now assume that we have a task $(\mathcal{S}, \mathcal{P})$ and a nested separation system $\mathcal{N}$ that is nested with $\mathcal{S}$, and which is oriented by every profile in $\mathcal{P}$ (like those obtained by applying one of $\{\mathsf{ext}, \mathsf{loc}, \mathsf{max}\}$ to $(\mathcal{S}, \mathcal{P})$). We shall define what it means to restrict a task $(\mathcal{S}, \mathcal{P})$ to an $\mathcal{N}$-profile X. Let $\mathcal{S}_X$ be the set of all separations $(A, B) \in \mathcal{S}$ that split X and let $\mathcal{P}_X$ be the set of all profiles in $\mathcal{P}$ to which X extends. The pair $(\mathcal{S}_X, \mathcal{P}_X)$ is the *restriction of* $(\mathcal{S}, \mathcal{P})$ *to* X, denoted by $(\mathcal{S}, \mathcal{P})_X$.

Lemma 4.13. *For every $\mathcal{N}$-profile X such that $\mathcal{P}_X$ is not empty, the restriction of $(\mathcal{S}, \mathcal{P})$ to X is a subtask of $(\mathcal{S}, \mathcal{P})$.*

Proof. Let X be an $\mathcal{N}$-profile such that $\mathcal{P}_X$ is not empty. In order to show that $(\mathcal{S}, \mathcal{P})_X$ is a subtask of $(\mathcal{S}, \mathcal{P})$ we have to show that $\mathcal{S}_X$ separates $\mathcal{P}_X$ well, and that $\mathcal{P}_X$ is $\mathcal{S}_X$-distinguishable. Since $\mathcal{S}$ is nested with $\mathcal{N}$ we have:

Every $\mathcal{P}'$-relevant separation in $\mathcal{S}$ splits X and is thus contained in $\mathcal{S}_X$. (4.8)

So since $\mathcal{S}$ distinguishes $\mathcal{P}'$ so does $\mathcal{S}_X$.

To show that $\mathcal{S}_X$ separates $\mathcal{P}_X$ well consider a crossing pair of separations $(A, B) \nparallel (C, D) \in \mathcal{S}_X$ and let P, P' be distinct profiles in $\mathcal{P}_X$ such that $(A, B), (C, D) \in P$ and $(B, A), (D, C) \in P'$. Since $\mathcal{S}$ separates $\mathcal{P}$ well, there is $(E, F) \in P \cap \mathcal{S}$ with $(A \cup C, B \cap D) \leq (E, F)$. Now by (P1) and the fact that P' orients $\mathcal{S}$ we have $(F, E) \in P'$, such that (E, F) distinguishes P from P'. By (4.8) this yields $(E, F) \in P \cap \mathcal{S}_X$. Hence, $\mathcal{S}_X$ separates $\mathcal{P}_X$ well. $\square$

A *strategy* is a map $\sigma : \mathbb{N} \to \{\mathsf{ext}, \mathsf{loc}, \mathsf{max}, \mathsf{ext}_{\mathcal{R}}, \mathsf{loc}_{\mathcal{R}}, \mathsf{max}_{\mathcal{R}}\}$ such that $\sigma(i) \in \{\mathsf{ext}_{\mathcal{R}}, \mathsf{loc}_{\mathcal{R}}, \mathsf{max}_{\mathcal{R}}\}$ for infinitely many $i \in \mathbb{N}$.

Theorem 4.14. *Let $(\mathcal{S}, \mathcal{P})$ be a task. Every strategy σ defines a canonical nested subsystem $\mathcal{N}_\sigma(\mathcal{S}, \mathcal{P}) \subseteq \mathcal{S}$ which distinguishes each pair of profiles in $\mathcal{P}$.*

Proof. Let r be the least integer such that $\sigma(r) \in \{\mathsf{ext}_\mathcal{R}, \mathsf{loc}_\mathcal{R}, \mathsf{max}_\mathcal{R}\}$ and let $p := |\mathcal{P}|$. We apply induction first on p and then on r.

If $p = 1$, then there is nothing to distinguish, such that $\mathcal{N}_\sigma = \emptyset$ has the required properties. So let us consider the case $p > 1$.

Let σ' be the strategy obtained from σ by 'dropping its first instruction', that is $\sigma'(i) = \sigma(i+1)$, for all $i \in \mathbb{N}$, and let r' be the least integer such that $\sigma'(r') \in \{\mathsf{ext}_\mathcal{R}, \mathsf{loc}_\mathcal{R}, \mathsf{max}_\mathcal{R}\}$. Consider the nested separation system $\mathcal{N}$ that is obtained from $(\mathcal{S}, \mathcal{P})$ by applying the first element of σ, that is $\mathcal{N} = \sigma(1)(\mathcal{S}, \mathcal{P})$.

If $r = 1$, we have $\sigma(1) \in \{\mathsf{ext}_\mathcal{R}, \mathsf{loc}_\mathcal{R}, \mathsf{max}_\mathcal{R}\}$. Let $(\mathcal{R}, \mathcal{P})$ be the reduction of the task $(\mathcal{S}, \mathcal{P})$. By definition of ext, loc and max, respectively, we have $\mathcal{N} \parallel \mathcal{R}$. Furthermore, due to (4.7) we have $\mathcal{N} \neq \emptyset$. Let X be an $\mathcal{N}$-profile such that $\mathcal{P}_X$ is not empty. Since $\mathcal{N}$ is nonempty it will distinguish some profiles in $\mathcal{P}$, such that $|\mathcal{P}_X| < |\mathcal{P}|$. So the induction hypothesis for p applied to σ' and $(\mathcal{R}, \mathcal{P})_X$ – which due to Lemma 4.13 is a subtask of $(\mathcal{R}, \mathcal{P})$ – yields a canonical nested subsystem $\mathcal{N}_{\sigma'}(\mathcal{R}, \mathcal{P})_X \subseteq \mathcal{R}_X$ which distinguishes all pairs of profiles in $\mathcal{P}_X$. By Lemma 4.12 we have $\mathcal{N}_{\sigma'}(\mathcal{R}, \mathcal{P})_X \parallel \mathcal{N}_{\sigma'}(\mathcal{R}, \mathcal{P})_Y$ for distinct $\mathcal{N}$-profiles X and Y. Let $\mathcal{X}$ denote the set of all those $\mathcal{N}$-profiles X for which $\mathcal{P}_X$ is not empty. Then

$$\mathcal{N}_\sigma(\mathcal{S}, \mathcal{P}) := \mathcal{N} \cup \bigcup_{X \in \mathcal{X}} \mathcal{N}_{\sigma'}(\mathcal{R}, \mathcal{P})_X$$

is a nested subsystem of $\mathcal{S}$ which distinguishes each pair of profiles in $\mathcal{P}$.

If $r > 1$, we have $\sigma(1) \notin \{\mathsf{ext}_\mathcal{R}, \mathsf{loc}_\mathcal{R}, \mathsf{max}_\mathcal{R}\}$. Hence we have $r' < r$ and we may apply the induction hypothesis for r to $(\mathcal{S}, \mathcal{P})_X$ and σ' for every $\mathcal{N}$-profile X for which $\mathcal{P}_X$ is not empty. With the same argument and using the same notation as in the previous case we obtain that

$$\mathcal{N}_\sigma(\mathcal{S}, \mathcal{P}) := \mathcal{N} \cup \bigcup_{X \in \mathcal{X}} \mathcal{N}_{\sigma'}(\mathcal{S}, \mathcal{P})_X$$

has the desired properties. $\square$

We now give an example to demonstrate that the use of different strategies may result in different nested separation systems (and therefore also to different tree-decompositions). Let Ext denote the strategy with $\mathsf{Ext}(i) = \mathsf{ext}_\mathcal{R}$ for all $i \in \mathbb{N}$, and let Loc and Max denote the analog strategies with $\mathsf{Loc}(i) = \mathsf{loc}_\mathcal{R}$ and $\mathsf{Max}(i) = \mathsf{max}_\mathcal{R}$, respectively. Notice that, even though (4.7) might have suggested the opposite, Example 4.1 shows that the separation systems obtained by Ext, Loc, or Max, respectively, may be pairwise incomparable under set inclusion; indeed no two of them are nested with each other.

Example 4.1. Let G be the 3-connected graph obtained from a (3×17)-grid by attaching on each 'short end' of the grid a K^4 and adding some edges as depicted in Figure 4.1. Let $\mathcal{S}$ be the set of all its 3-separations and $\mathcal{P}$ the set of $\mathcal{S}$-block-profiles that correspond to the 4-blocks of G. It is straightforward to verify that $(\mathcal{S}, \mathcal{P})$ is a task. The gray bars in Figures 4.1 (a) to (c) highlight the separators of the separations in (a) $\mathcal{N}_{\mathsf{Ext}}(\mathcal{S}, \mathcal{P})$, (b) $\mathcal{N}_{\mathsf{Loc}}(\mathcal{S}, \mathcal{P})$, and (c) $\mathcal{N}_{\mathsf{Max}}(\mathcal{S}, \mathcal{P})$, respectively.

In case (a), at each step we split off the two outmost 4-blocks in the middle block, until we end up with a single 4-block in the middle after four steps.

Figure 4.1: Distinguishing the 4-blocks of a 3-connected graph using either:

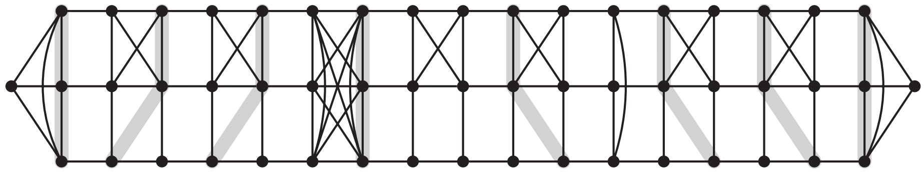

(a) extremal separations (Ext),

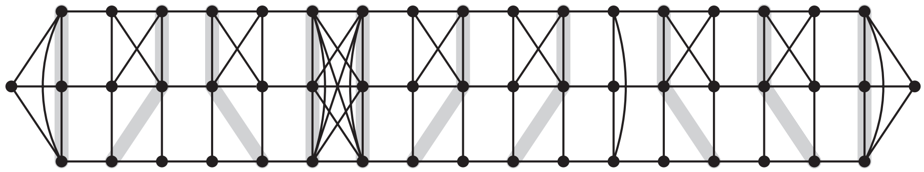

(b) locally-extremal separations (Loc),

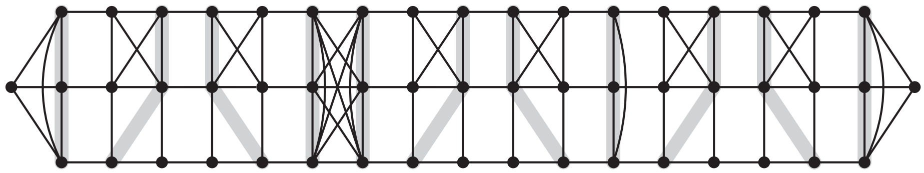

(c) or all possible separations (Max).

In case (b) we split off the two outmost K^4's as well as the K^6 in the first step by means of the 'straight separators' of Figure 4.1 (b), producing two blocks that contain more than one 4-block. In the second step we distinguish the two K^4's in the left block and split off the two outer K^4's in the right block. In the third step we finally distinguish the two K^4's which still lie in a common block.

In the first step of case (c) we choose, in addition to those separations chosen in the first step of (b), the two 3-separations whose common separator forms a maximal clique, highlighted by the additional straight bar in Figure 4.1 (c). This produces three blocks containing two K^4's, which we distinguish in the second and last step of (c).

4.4 Distinguishing the k-profiles of a connectivity system

Let V be a finite set. A *separation lattice* over V is a symmetric set $\mathcal{C}$ of separations of V such that for all pairs $(A, B), (C, D) \in \mathcal{C}$ we also have $(A \cup C, B \cap D) \in \mathcal{C}$. The set V is the *groundset* of $\mathcal{C}$ and we will later refer to it by $V(\mathcal{C})$. Notice that by symmetry, a separation lattice will contain all corner-separations of all pairs of separations it contains. A *connectivity system* is a separation lattice $\mathcal{C}$ together with a function $\lambda : \mathcal{C} \to \mathbb{N}$ which satisfies:

(C1) $\lambda(A, B) = \lambda(B, A)$;

(C2) $\lambda(A, B) + \lambda(C, D) \geq \lambda(A \cup C, B \cap D) + \lambda(A \cap C, B \cup D)$.

Such a function λ is called a *connectivity function* on $\mathcal{C}$. Property (C1) is referred to as *symmetry* while property (C2) is referred to as *sub-modularity*. Notice that our definition of a connectivity system is 'essentially' the same as the one given by Geelen, Gerards and Whittle in [9]. One way in which we shall make use of this notion of a connectivity system is in order to proof some results about tangles in matroids (see Section 5.3). So our first example of a connectivity system will be obtained from a matroid M:

Example 4.2. The separations of a matroid M on a ground-set E with rank function r, together with the connectivity function λ sending a separation (X, Y) of M to $\lambda(X, Y) = r(X) + r(Y) - r(M) + 1$ is a connectivity system over E.

For a proof of Example 4.2 see for example Oxley [13]. It is a well-known fact that a graph G gives rise to a connectivity system (in the sense of [9]) over the set of edges of G. With our slightly modified definition of a connectivity system we can also associate G wit a connectivity system over its set of vertices:

Example 4.3. Given a graph G let $\mathcal{C}_G$ be the set of separations of G together with the function $\lambda : (A, B) \mapsto |A \cap B|$, which by Lemma 3.1 is a connectivity function on $\mathcal{C}_G$. In fact, for the statement of Lemma 3.1 we took it for granted, without proving it, that the corner-separations of a pair of separations of G are again separations of G. However, the proof for this fact is very simple and we will not go into details here either. We conclude that $\mathcal{C}_G$ is indeed a connectivity system over the set of vertices of G.

For the remainder of this section let $\mathcal{C}$ be a connectivity system. Motivated by Examples 4.2 and 4.3 we call the value $\lambda(A, B)$ of a separation $(A, B) \in \mathcal{C}$ the *order* of (A, B), and as for graphs and matroids we call a separation of order k a k-separation. A profile $P \subseteq \mathcal{C}$ is called a profile *of* the connectivity system $\mathcal{C}$. A profile P of a connectivity system $\mathcal{C}$ is a *k-profile* of $\mathcal{C}$ if it has, in addition to (P1) and (P2), the following property:

Every separation in P has order $< k$, and for every separation (A, B) of order $< k$ exactly one of (A, B) and (B, A) lies in P. (4.9)

A profile P that is a k-profile, for some $k \in \mathbb{N}$, is called a *regular* profile. Given a regular profile P the unique integer k such that P is a k-profile is the *order* of P, denoted by $\mathrm{ord}(P)$.

Let $\mathcal{S}_{<k}$ denote the set of all separations of order $< k$ of $\mathcal{C}$. A profile P is called k-complete if $P \cap \mathcal{S}_{<k}$ is a k-profile. That is, a profile is k-complete if and only if it contains a k-profile. Given two profiles P_1 and P_2 let $\kappa(P_1, P_2)$ be the smallest integer k such that both P_1 and P_2 are k-complete and we have $P_1 \cap \mathcal{S}_{<k} \neq P_2 \cap \mathcal{S}_{<k}$; put $\kappa(P_1, P_2) = \infty$, if no such k exists. Two profiles P_1, P_2 are called *k-distinguishable* if $\kappa(P_1, P_2) \leq k$. It is easy to see that if P_1 and P_2 are k-distinguishable then there is a separation of order $< k$ that distinguishes P_1 and P_2. The converse, however, need not hold. If all we know is that (A, B) is a separation of order $< k$ that distinguishes two profiles P_1 and P_2, then we do not know whether or not P_1 and P_2 are both k-complete, which is a necessary condition for P_1 and P_2 in order to be k-distinguishable. But two k-complete profiles P_1 and P_2 are k-distinguishable if and only if there is a separation of order $< k$ that distinguishes them. In particular we have:

If $\kappa(P_1, P_2) = k$ then the minimum order of a separation that distinguishes P_1 and P_2 is $k - 1$. (4.10)

A separation (A, B) that distinguishes two profiles P_1 and P_2 does so *efficiently* if we have $\lambda(A, B) = \kappa(P_1, P_2) - 1$. A set of separations distinguishes a pair of profiles *efficiently* if it contains a separation that does so.

A regular profile P is called *K-robust* if for every separation $(C, D) \in P$ and every separation (A, B) of order $< K$ such that

$$\lambda(A \cup C, B \cap D) < \mathrm{ord}(P) - 1 > \lambda(B \cup C, A \cap D)$$

we have either $(A \cup C, B \cap D) \in P$ or $(B \cup C, A \cap D) \in P$. (Note that due to (2.5) and (P1) there cannot be both in P.) By (P2) we have:

Every k-profile is k-robust. (4.11)

Furthermore we obtain the following lemma:

Lemma 4.15. *Let P be a K-robust profile (which is regular). Then every regular profile $P' \subseteq P$ is K'-robust for all $K' \leq K$.*

Proof. Let P be a K-robust profile and consider a regular profile $P' \subseteq P$ and let $K' \leq K$. Suppose that P' is not K'-robust. Then there is a separation (A, B) of order $< K'$ and a separation $(C, D) \in P'$ such that both $(A \cup C, B \cap D)$ and $(B \cup C, A \cap D)$ have order less than $\mathrm{ord}(P') - 1$ but neither of these two separations is contained in P'. By (4.9) this yields that both $(B \cap D, A \cup C)$ and $(A \cap D, B \cup C)$ are contained in P'. Since $P' \subseteq P$ we have $\mathrm{ord}(P') \leq \mathrm{ord}(P)$, and since $K' \leq K$ we have $\lambda(A, B) < K$. Hence, (A, B) and (C, D) witness that P is not K-robust, contradicting our assumption. □

A pair P_1, P_2 of profiles of $\mathcal{C}$ is *K-robustly* k-distinguishable if it is k-distinguishable, that is if $\kappa := \kappa(P_1, P_2) \leq k$, and if both $P_1 \cap \mathcal{S}_\kappa$ and $P_2 \cap \mathcal{S}_\kappa$ are K-robust. We are now ready to prove the main result of this thesis.

Theorem 4.16. *Let $\mathcal{C}$ be a connectivity system and let $K > 0$ be an integer. Then for every set $\mathcal{P}$ of profiles of $\mathcal{C}$ there is a nested separation system $\mathcal{N}(\mathcal{P})$ that efficiently distinguishes each K-robustly K-distinguishable pair of profiles in $\mathcal{P}$. Furthermore, every separation in $\mathcal{N}(\mathcal{P})$ separates such a pair efficiently.*

Proof. The proof follows essentially the same approach as the proof of Theorem 3.6. For $0 < k \leq K$ let $\mathcal{P}_k$ denote the set of all K-robust k-profiles that are a subset of a profile in $\mathcal{P}$. Recursively for all integers $0 < k \leq K$ we shall construct a sequence of nested separation systems $\mathcal{N}_k$ with the following properties:

(i) $\mathcal{N}_{k-1} \subseteq \mathcal{N}_k$ (put $\mathcal{N}_0 := \emptyset$);

(ii) every separation in $\mathcal{N}_k \setminus \mathcal{N}_{k-1}$ has order $k - 1$;

(iii) $\mathcal{N}_k$ distinguishes each pair of profiles in $\mathcal{P}_k$;

(iv) every separation in $\mathcal{N}_k \setminus \mathcal{N}_{k-1}$ distinguishes a pair of profiles in $\mathcal{P}_k$ that are not distinguished by $\mathcal{N}_{k-1}$;

Let us first show that $\mathcal{N}(\mathcal{P}) := \mathcal{N}_K$ would then satisfy the assertions of the theorem. Consider a K-robustly K-distinguishable pair P_1, P_2 of profiles in $\mathcal{P}$, and let $m := \kappa(P_1, P_2)$. Then we have $m \leq K$ and P_1 and P_2 each contain a K-robust m-profile P_1' and P_2', respectively, such that $P_1' \neq P_2'$. Hence, P_1' and P_2' are distinct profiles in $\mathcal{P}_m$. By property (iii) we know that $\mathcal{N}_m$ distinguishes P_1' and P_2' and therefore also P_1 and P_2, and so does $\mathcal{N}_K$ by property (i), it does so efficiently by definition of m and property (ii). Hence, $\mathcal{N}_K$ satisfies the first assertion of the theorem. The second follows easily from property (iv).

So let us construct the nested separation systems $\mathcal{N}_k$. Let $k > 0$ be given, and assume inductively that we already have nested separation systems $\mathcal{N}_{k'}$ satisfying (i) to (iv) for all $0 < k' < k$. For $k = 1$ we just have $\mathcal{N}_0 = \emptyset$.

If $\mathcal{N}_{k-1}$ already distinguishes $\mathcal{P}_k$, then $\mathcal{N}_k := \mathcal{N}_{k-1}$ will obviously satisfy (i) to (iv). If not, then there are profiles $P_1, P_2 \in \mathcal{P}_k$ that are not distinguished by $\mathcal{N}_{k-1}$. By definition of $\mathcal{P}_k$ both P_1 and P_2 are K-robust, so Lemma 4.15 implies that also $Q_i := P_i \cap \mathcal{S}_{k-1}$ is a K-robust profile, for $i \in \{1, 2\}$, which implies $Q_i \in \mathcal{P}_{k-1}$. By the induction hypothesis (iii), $\mathcal{N}_{k-1}$ distinguishes each pair of profiles in $\mathcal{P}_{k-1}$, hence Q_1 and Q_2 cannot be distinct. Let us put $Q := Q_1 = Q_2$ and let $\mathcal{P}_Q := \{P \in \mathcal{P}_k \mid Q \subseteq P\}$ (such that $P_1, P_2 \in \mathcal{P}_Q$). Furthermore let $\mathcal{S}_Q$ denote the set of all separations of order $k - 1$ that split $Q \cap \mathcal{N}_{k-1}$. We shall show that $(\mathcal{S}_Q, \mathcal{P}_Q)$ is a task. First we show:

($*$) *$\mathcal{S}_Q$ distinguishes $\mathcal{P}_Q$.*

Let P_1 and P_2 be arbitrary profiles in $\mathcal{P}_Q$. Since both contain the same $(k-1)$-profile Q we have $\kappa(P_1, P_2) = k$. So there is a $(k-1)$-separation which distinguishes P_1 from P_2. Let (A, B) be such a separation with $(A, B) \in P_1$ and $(B, A) \in P_2$ and which in addition is nested with as many separations in $\mathcal{N}_{k-1}$ as possible. If we show that (A, B) is in fact nested with $\mathcal{N}_{k-1}$ then it clearly splits $Q \cap \mathcal{N}_{k-1}$ (witnessed by P_1 and P_2) and must therefore be in $\mathcal{S}_Q$.

Suppose that (A, B) is not nested with $\mathcal{N}_{k-1}$. Then there is $(C, D) \in \mathcal{N}_{k-1}$ such that (A, B) and (C, D) cross. Let $\ell < k$ be such that $(C, D) \in \mathcal{N}_\ell \setminus \mathcal{N}_{\ell-1}$. By (ii) we know that (C, D) has order $\ell - 1$ and by (iv) there are profiles $R_1, R_2 \in \mathcal{P}_\ell$ with $(C, D) \in R_1$ and $(D, C) \in R_2$. Furthermore, by (iii) and (iv) together, no separation of order less than $\ell - 1$ distinguishes R_1 from R_2. Since (C, D) does not distinguish P_1 from P_2 we may assume $(D, C) \in P_1 \cap P_2$.

Assume that the corner-separation $(E, F) := (A \cap D, B \cap C)$ has order less than k. Since P_1 is a k-profile we then have $(F, E) \notin P_1$ due to (P2), which implies $(E, F) \in P_1$ due to (4.9). On the other hand we have $(E, F) \notin P_2$ due to (P1), and hence $(F, E) \in P_2$ by (4.9). That is, (E, F) distinguishes P_1 from P_2. Since we have $\kappa(P_1, P_2) = k$ this implies that (E, F) is a $(k-1)$-separation. But then (E, F) contradicts the choice of (A, B): by Lemma 2.2 (E, F) is nested with every separation of $\mathcal{N}_{k-1}$ with which (A, B) is nested, and (E, F) is also nested with (C, D), which (A, B) crosses.

Hence $(E, F) = (A \cap D, B \cap C)$ has order at least k, which by sub-modularity and symmetry of the connectivity function implies that the order of the corner-separation corresponding to the corner opposite to $B \cap C$ has order strictly less than $\ell - 1$. The symmetric argument, changing the roles of P_1 and P_2, shows that the corner-separation corresponding to the corner opposite to $A \cap C$ has order strictly less than $\ell - 1$. That is we have

$$\lambda(A \cup C, B \cap D) < \ell - 1 > \lambda(B \cup C, A \cap D).$$

By assumption we have $(C,D) \in R_1$. Since R_1 is a K-robust ℓ-profile, it must contain one of these separations, we may assume $(A \cup C, B \cap D) \in R_1$. But then $(A \cup C, B \cap D)$ distinguishes R_1 and R_2, since we have $(B \cap D, A \cup C) \in R_2$ due to (P1) and (4.9). This contradicts $\kappa(R_1, R_2) = \ell$ and therefore finishes the proof of $(*)$.

Next we have to show:

$(**)$ *$\mathcal{S}_Q$ separates $\mathcal{P}_Q$ well.*

Consider two crossing separations $(A,B),(C,D) \in \mathcal{S}_Q$ and two profiles $P_1, P_2 \in \mathcal{P}_Q$ such that $(A,B),(C,D) \in P_1$ and $(B,A),(D,C) \in P_2$. Suppose that the corner-separation $(A \cup C, B \cap D)$ has order at least k. Then by sub-modularity we have $\lambda(B \cup D, A \cap C) < k-1$. Then by the profile properties and (4.9) we have $(B \cup D, A \cap C) \in P_2$ and $(A \cap C, B \cup D) \in P_1$, contradicting $\kappa(P_1, P_2) = k$.

Hence the order of $(A \cup C, B \cap D)$ is less than k. So with the same argument as above it distinguishes P_1 from P_2, and therefore has order $k-1$. By Lemma 2.2 $(A \cup C, B \cap D)$ is nested with every separation with which both (A,B) and (C,D) are nested. Hence, $(A \cup C, B \cap D)$ is nested with $\mathcal{N}_{k-1}$ and is therefore contained in $\mathcal{S}_Q$, which finishes the proof of $(**)$.

So we know that $(\mathcal{S}_Q, \mathcal{P}_Q)$ is a task. By Theorem 4.7 there is a $\mathcal{P}_Q$-relevant nested separation system $\mathcal{N}(\mathcal{S}_Q, \mathcal{P}_Q)$ that distinguishes each pair of profiles in $\mathcal{P}_Q$. The same argument applies to all profiles $Q' \in \mathcal{P}_{k-1}$ that are contained in more than one profile of $\mathcal{P}_k$; let $\mathcal{Q}$ denote the subset of $\mathcal{P}_{k-1}$ of all those profiles. Let us construct $\mathcal{N}_k$ from $\mathcal{N}_{k-1}$ as follows:

$$\mathcal{N}_k := \mathcal{N}_{k-1} \cup \bigcup_{Q \in \mathcal{Q}} \mathcal{N}(\mathcal{S}_Q, \mathcal{P}_Q).$$

We have to check that $\mathcal{N}_k$ satisfies (i) to (iv). Assertions (i), (ii) and (iv) are obvious due to the construction of $\mathcal{N}_k$. To verify (iii) consider distinct profiles $P_1, P_2 \in \mathcal{P}_k$. For $i \in \{1,2\}$ let $Q_i := P_i \cap \mathcal{S}_{k-1}$. By Lemma 4.15 we have $Q_i \in \mathcal{P}_{k-1}$. If $Q_1 \neq Q_2$ then $\mathcal{N}_{k-1}$ distinguishes Q_1 from Q_2, due to the induction hypothesis. Hence $\mathcal{N}_{k-1}$ also distinguishes P_1 from P_2, and so does $\mathcal{N}_k$, which is a superset of $\mathcal{N}_{k-1}$. On the other hand, if we have $Q := Q_1 = Q_2$, then Q is contained in $\mathcal{Q}$. Hence, $\mathcal{N}_k$ contains $\mathcal{N}(\mathcal{S}_Q, \mathcal{P}_Q)$ as a subset, which distinguishes P_1 from P_2. This finishes the construction of $\mathcal{N}_k$. □

A permutation α of $V(\mathcal{C})$ is an *automorphism* of the connectivity system $\mathcal{C}$ if, for every $(A,B) \in \mathcal{C}$, we have $(A,B)^\alpha \in \mathcal{C}$ and $\lambda(A,B) = \lambda((A,B)^\alpha)$. As usual we denote the group of automorphisms of $\mathcal{C}$ by $\mathrm{Aut}(\mathcal{C})$. From Corollary 4.8 and the way we we construct the separation system $\mathcal{N}(\mathcal{P})$ we easily obtain the following corollary.

Corollary 4.17. *For every automorphism α of $\mathcal{C}$ we have $\mathcal{N}(\mathcal{P}^\alpha) = \mathcal{N}(\mathcal{P})^\alpha$.*

While our proof of Theorem 4.16 was based on Theorem 4.7, we could instead have used its 'refined' version Theorem 4.14 to gain some flexibility. We have to take care, however, that every separation that we chose when solving some task is relevant with respect to the corresponding set of profiles. We do so by choosing at every 'step k' within the proof of Theorem 4.16, that is for every $0 < k \leq K$, a strategy σ_k with $\sigma_k(i) \in \{\mathsf{ext}_\mathcal{R}, \mathsf{loc}_\mathcal{R}, \mathsf{max}_\mathcal{R}\}$ for all $i \in \mathbb{N}$; let us call a K-tuple $\Sigma := (\sigma_1, \cdots, \sigma_K)$ of such strategies a *K-strategy*. So we obtain the following corollary from the proof of Theorem 4.16. (A formal proof can be found in [2].)

Corollary 4.18. *Every K-strategy Σ induces for every set $\mathcal{P}$ of profiles a canonical nested separation system $\mathcal{N}_\Sigma(\mathcal{P})$ that efficiently distinguishes each K-robustly K-distinguishable pair of profiles in $\mathcal{P}$. Furthermore, every separation in $\mathcal{N}_\Sigma(\mathcal{P})$ separates such a pair efficiently.* □

A *tree-decomposition of $\mathcal{C}$* is a tree-decomposition of $V(\mathcal{C})$ such that the nested separation system it induces is a subset of $\mathcal{C}$. The *adhesion* of a tree-decomposition of $\mathcal{C}$ is the maximum order of a separation it induces. As in Section 3.3, we now turn the nested separation system obtained from Theorem 4.16 into a sequence of tree-decompositions of $\mathcal{C}$.

Theorem 4.19. *Let $\Gamma \subseteq \mathrm{Aut}(\mathcal{C})$ be a group of automorphisms of a connectivity system $\mathcal{C}$ and let $\mathcal{P}$ be a Γ-invariant set of profiles of $\mathcal{C}$. Then for every positive integer K there is a sequence $(\mathcal{T}_k, \mathcal{V}_k)_{k \le K}$ of tree-decompositions of $\mathcal{C}$ such that, for all $0 < k \le K$,*

(i) *every k-profile in $\mathcal{P}$ inhabits a unique part of $(\mathcal{T}_k, \mathcal{V}_k)$;*

(ii) *distinct K-robust k-profiles in $\mathcal{P}$ inhabit different parts of $(\mathcal{T}_k, \mathcal{V}_k)$;*

(iii) *$(\mathcal{T}_k, \mathcal{V}_k)$ has adhesion less than k;*

(iv) *if $k > 1$ then $(\mathcal{T}_{k-1}, \mathcal{V}_{k-1}) \preccurlyeq (\mathcal{T}_k, \mathcal{V}_k)$;*

(v) *Γ acts on $\mathcal{T}_k$ as a group of automorphisms.*

Proof. Let $\mathcal{N} := \mathcal{N}(\mathcal{P})$ be the nested separation system obtained for $\mathcal{P}$ and K due to Theorem 4.16. For every $0 < k < K$ let $\mathcal{N}_k := \mathcal{N} \cap \mathcal{S}_{<k}$ and let $(\mathcal{T}_k, \mathcal{V}_k)$ be the tree-decomposition of $\mathcal{C}$ obtained from $\mathcal{N}_k$. By Corollaries 2.9 and 4.17 the group Γ will act on each $\mathcal{T}_k$ as a group of automorphisms, which proves (v). Statements (iii) and (iv) are due to Theorem 2.17 (iii) and (iv) and statement (i) is due to Corollary 4.4.

For a proof of (ii) consider distinct K-robust k-profiles P_1, P_2. They are clearly K-robustly k-distinguishable and therefore also K-robustly K-distinguishable. So by Theorem 4.16 there is a separation $(A, B) \in \mathcal{N}$ that efficiently distinguishes P_1 from P_2. Hence the order of (A, B) is $\kappa(P_1, P_2) - 1 < k$, which implies $(A, B) \in \mathcal{N}_k$. Thus we have $P_1 \cap \mathcal{N}_k \neq P_2 \cap \mathcal{N}_k$, which shows that P_1 and P_2 inhabit different parts of $(\mathcal{T}_k, \mathcal{V}_k)$. □

Let us say that a tree-decomposition *(efficiently) distinguishes* a pair of profiles if the separation system it induces does so. As another application of Theorem 4.16 we obtain the following result.

Theorem 4.20. *Let $\Gamma \subseteq \mathrm{Aut}(\mathcal{C})$ be a group of automorphisms of a connectivity system $\mathcal{C}$, let $K > 0$ be an integer and let $\mathcal{P}$ be a Γ-invariant set of pairwise K-robustly K-distinguishable profiles of $\mathcal{C}$. Then there is a tree-decomposition $(\mathcal{T}, \mathcal{V})$ of $\mathcal{C}$ such that*

(i) *every profile in $\mathcal{P}$ inhabits a unique part of $(\mathcal{T}, \mathcal{V})$;*

(ii) *$(\mathcal{T}, \mathcal{V})$ efficiently distinguishes each pair of profiles in $\mathcal{P}$;*

(iii) $(\mathcal{T}, \mathcal{V})$ *has adhesion less than* K*;*

(iv) Γ *acts on* $\mathcal{T}$ *as a group of automorphisms.*

Proof. Let $\mathcal{N} := \mathcal{N}(\mathcal{P})$ be the nested separation system obtained for $\mathcal{P}$ and K due to Theorem 4.16 and let $(\mathcal{T}, \mathcal{V})$ be the tree-decomposition of $\mathcal{C}$ obtained from $\mathcal{N}$. Statements (ii) to (iv) can be proved with the same arguments found in the proof of Theorem 4.19.

For a proof of (i) consider a profile $P \in \mathcal{P}$ and let us define

$$k(P) := \max\{\kappa(P, P') \mid P \neq P' \in \mathcal{P}\}. \tag{4.12}$$

Then P is $k(P)$-complete and therefore orients $\mathcal{N}' := \mathcal{N} \cap \mathcal{S}_{k(P)}$. Let $P' := P \cap \mathcal{N}'$ be the $\mathcal{N}'$-profile contained in P. Consider an arbitrary separation (C, D) that is contained in $\mathcal{N}$ but not in $\mathcal{N}'$. By Theorem 4.16 there are profiles $P_1, P_2 \in \mathcal{P}$ so that (C, D) efficiently distinguishes them. That is we have

$$\kappa(P_1, P_2) > \lambda(C, D) \geq k(P). \tag{4.13}$$

Hence, by (4.12) and (4.13) we know that P, P_1, P_2 are pairwise distinct $k(P)$-complete profiles. We also obtain from (4.13), that P_1 and P_2 contain the same $\mathcal{N}'$-profile $Q' := P_1 \cap \mathcal{N}' = P_2 \cap \mathcal{N}'$. Therefore (C, D) splits Q', which by Lemma 4.10 implies that (C, D) does not split P'. Since (C, D) was chosen arbitrary in $\mathcal{N} \setminus \mathcal{N}'$ Lemma 4.11 implies that P' extends to a unique $\mathcal{N}$-preference. Hence, P inhabits a unique part of $(\mathcal{T}, \mathcal{V})$. $\square$

The most natural application of Theorem 4.19 will be to construct a tree-decomposition $(\mathcal{T}, \mathcal{V})$ which efficiently distinguishes a set $\mathcal{P}$ of k-profiles, for some integer $k > 0$. Let us call a part of $(\mathcal{T}, \mathcal{V})$ *essential* if it is inhabited by a k-profile in $\mathcal{P}$, and *inessential* otherwise.

By Corollary 4.18 we have different strategies available to construct the underlying nested separation system from which we obtain $(\mathcal{T}, \mathcal{V})$. In that way we can reduce the number of inessential parts of our tree-decomposition to a minimum: we obtain bounds on the number of separations in the underlying nested separation system $\mathcal{N}(\mathcal{P})$ in terms of the number of k-profiles in $\mathcal{P}$. Since each edge of the corresponding decomposition tree corresponds to a pair $\{(A, B), (B, A)\} \subseteq \mathcal{N}$ (see Section 2.2), and the number of essential parts is exactly the number of k-profiles in $\mathcal{P}$, we can easily obtain the number of inessential parts from these bounds. The most obvious method to keep the amount of chosen separations low, is to consider only relevant separations, i.e. those that actually distinguish two k-profiles. Then, (4.7) may suggest to use only $\mathsf{ext}_{\mathcal{R}}$ – that is, to choose the strategy Ext. Example 4.4 will show, however, that there are graphs for which the use of Loc yields better results.

Given a set $\mathcal{P}$ of k-profiles let $p := |\mathcal{P}|$. It is clear that any nested separation system that distinguishes all pairs of distinct profiles in $\mathcal{P}$, must contain at least $2(p-1)$ separations. We obtain the following upper bounds.

Lemma 4.21. *For every task* $(\mathcal{S}, \mathcal{P})$ *we have:*

$$|\mathcal{N}_{\mathsf{Ext}}(\mathcal{S}, \mathcal{P})| \leq 2p, \text{ and} \tag{4.14}$$

$$|\mathcal{N}_{\mathsf{Loc}}(\mathcal{S}, \mathcal{P})| \leq 4(p-1). \tag{4.15}$$

Proof. Let us prove (4.14) first. Let (A, B) be a separation in $\mathcal{N}_{\mathsf{Ext}}(\mathcal{S}, \mathcal{P})$. Then either (A, B) or (B, A) was extremal in some reduced subtask $(\mathcal{S}', \mathcal{P}')$; we may assume this was (A, B). Then there is a unique profile $P_{(A,B)} \in \mathcal{P}'$ with $(A, B) \in P_{(A,B)}$ and for every $(A, B) \neq (C, D) \in \mathsf{ext}(\mathcal{S}', \mathcal{P}')$ we have $(C, D) \notin P_{(A,B)}$. Furthermore, $P_{(A,B)}$ gets distinguished by (A, B) from every other profile in $\mathcal{P}'$. Hence, every profile in $\mathcal{P}$ justifies the choice of at most 2 separations, which shows (4.14).

Now let $(\mathcal{T}, \mathcal{V})$ be the tree-decomposition obtained from $\mathcal{N}_{\mathsf{Loc}}(\mathcal{S}, \mathcal{P})$. We shall show the following:

$$\textit{If } \{t_1, t_2\} \textit{ is an edge of } \mathcal{T}\textit{, then either } V_{t_1} \textit{ or } V_{t_2} \textit{ is essential.} \tag{4.16}$$

Suppose this is false and let $e = \{t_1, t_2\}$ be a witness for this fact. Let $\mathcal{T}_1$ and $\mathcal{T}_2$ denote the components of $\mathcal{T} - e$ which contain t_1 and t_2, respectively. Then e corresponds to a pair $\{(A, B), (B, A)\}$ of separations in $\mathcal{S}$, one of which must have been locally-extremal in a certain reduced subtask $(\mathcal{S}', \mathcal{P}')$; we may assume this was (A, B). That is, there is a profile $P \in \mathcal{P}'$ which is well-separated in $(\mathcal{S}', \mathcal{P}')$ and such that (A, B) is $\leq$-maximal in $P \cap \mathcal{S}'$. Since $(\mathcal{T}, \mathcal{V})$ has adhesion less than k, P must inhabit a part V_t, and by assumption we have $t \notin \{t_1, t_2\}$. We may assume that t lies in $\mathcal{T}_1$, such that $V_{t'} \subseteq B$, for all $t' \in \mathcal{T}_1$. As $(\mathcal{S}', \mathcal{P}')$ is reduced, (A, B) is $\mathcal{P}'$-relevant, so there must be another profile $P \neq P' \in \mathcal{P}'$ with $(B, A) \in P'$. Let e' be the edge incident with t on the unique t_1–t-path in $\mathcal{T}$. Then e' corresponds to a pair $\{(C, D), (D, C)\} \subseteq \mathcal{S}$. We may assume $V_t \subseteq D$, such that $(A, B) \leq (C, D)$. But then we have $(C, D) \in P$, since P inhabits V_t, and $(D, C) \in P'$, due to (P1). This implies that (C, D) is $\mathcal{P}'$-relevant, and therefore contained in $\mathcal{S}'$ (cp. proof of Lemma 4.13). So (A, B) was not $\leq$-maximal in $P \cap \mathcal{S}'$, a contradiction. This shows (4.16).

Since we only choose $\mathcal{P}$-relevant separations, all the parts corresponding to a leaf of $\mathcal{T}$ must be essential. That is, every node of $\mathcal{T}$ such that V_t is inessential, has at least 2 neighbours, all of which, by (4.16), correspond to essential parts. Let us say that *suppressing* a vertex of a tree means to delete it and to add an edge from one of its neigbours to all the other of its neighbours. Then by suppressing all the nodes of $\mathcal{T}$ that correspond to inessential parts, we obtain a tree $\mathcal{T}'$ with exactly p vertices, which has at least half the number of edges as $\mathcal{T}$. That is, $\mathcal{T}$ has at most $2(p-1)$ edges. Since each edge of $\mathcal{T}$ corresponds to a pair of separations in $\mathcal{N}_{\mathsf{Loc}}(\mathcal{S}, \mathcal{P})$, we obtain (4.15). □

Now we come to prove the before-mentioned bounds on the number of inessential parts. Let Ext^k denote the k-strategy all whose entries are Ext and let Loc^k denote the k-strategy which only uses Loc. Then we have:

Theorem 4.22. *For every set $\mathcal{P}$ of k-profiles we have*

$$|\mathcal{N}_{\mathsf{Ext}^k}(\mathcal{P})| \leq 4(p-1) \geq |\mathcal{N}_{\mathsf{Loc}^k}(\mathcal{P})|.$$

Proof. We define a rooted tree (T, r) that represents the iterative approach of the proof of Theorem 4.16 as follows: for $0 < \ell \leq k$ let $\mathcal{P}_\ell$ denote the set of those ℓ-profiles that are a subset of a profile in $\mathcal{P}$ and let the vertex set of T be given by

$$V(T) := \{\emptyset\} \cup \bigcup_{0 < \ell \leq k} \mathcal{P}_\ell,$$

we set $r = \emptyset$ as the root, and we add an edge between v and w if one is the predecessor of the other with repsect to set inclusion. Then the set of leaves, where the root is not considered as a leaf, is precisely the set of k-profiles $\mathcal{P}$. Let us call the non-leaves of T (including the root) its *internal* vertices.

Let v be an internal vertex of T. Then all children of v are ℓ_v-profiles for some $0 < \ell_v \leq k$. Now every task $(\mathcal{S}', \mathcal{P}')$ that we solve during the iteration is represented by a vertex v' such that $\mathcal{P}'$ is precisely the set of children of v'. And vice versa: there is a one-to-one correspondence between the internal vertices of T and the tasks that we have to solve. Let $c(v)$ denote the number of children of an internal vertex v. Since we only consider Ext or Loc, respectively, to solve the tasks that occur, we will only produce non-empty separation systems for tasks $(\mathcal{S}', \mathcal{P}')$ that correspond to a vertex v' with $c(v') = |\mathcal{P}'| \geq 2$. So for the aim of this proof, it is reasonable to suppress in T all internal vertices with less than two children, while suppressing the root turns its child into the new root. This results in a rooted tree (T', r') whose internal vertices have degree at least 3, except for the root, which has degree at least 2. Let i denote the number of internal vertices in T'. Since the number of leaves of T' is exactly p, we have at most $(p-1)$ internal vertices, that is we have $i \leq p-1$.

Now consider the construction of $\mathcal{N}_{\mathsf{Ext}^k}(\mathcal{P})$. By (4.14), each internal vertex v of T' contributes at most $2c(v)$ separations. So there are at most twice the number of edges of T' many separations in $\mathcal{N}_{\mathsf{Ext}^k}(\mathcal{P})$. Hence, we have

$$|\mathcal{N}_{\mathsf{Ext}^k}(\mathcal{P})| \leq 2|E(T')| = 2(p+i-1) \leq 4(p-1).$$

During the construction of $\mathcal{N}_{\mathsf{Loc}^k}(\mathcal{P})$, each internal vertex v of T' contributes at most $4(c(v)-1)$ separations, due to (4.15). So if we say that each edge of T' contributes four separations, then we add a surplus of four separations per internal vertex. That is we have:

$$|\mathcal{N}_{\mathsf{Loc}^k}(\mathcal{P})| \leq 4|E(T')| - 4i = 4(p+i-1) - 4i = 4(p-1).$$

This finishes the proof. $\square$

It is remarkable that Lemma 4.21 establishes a significantly better bound for Ext than for Loc when applied to a single task, while Theorem 4.22 gives the same upper bound for both Ext^k and Loc^k, which correspond to repetitive applications of Ext and Loc, respectively. But still, the lower and upper bound on both Ext^k and Loc^k are sharp. Examples where the upper bound is attended are quite easy to construct. Here, we want to give an example where Loc yields the best possible result of $2(p-1)$, while Ext does not.

Example 4.4. A modification of the graph of Example 4.1 yields a 3-connected graph with four 4-blocks, as shown in Figure 4.2. As in Figure 4.1 the gray bars represent separators of chosen separations. Here Ext chooses all corresponding separations, first the outer ones, then in a second step the inner ones. On the other hand, Loc will choose all separations corresponding to the 'straight' separators at the first step, such that no second step is needed. Therefore Ext chooses two separations more than Loc.

62

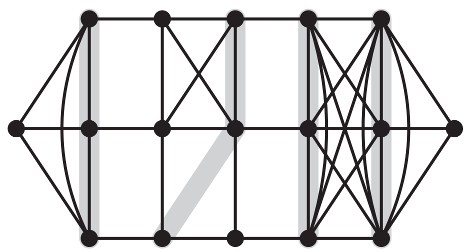

Figure 4.2: A graph where Loc chooses fewer separations than Ext.

Chapter 5

The k-profiles of graphs and matroids

In this chapter we apply the results of the previous chapters to finite graphs and matroids. In Section 5.1 we define k-profiles of graphs and show that this notion generalises both the notion of a k-block and the notion of a tangle of order k. We also give a characterisation of the k-profiles of a graph using the notion of a k-haven. In Section 5.2 we prove the main results for graphs that were stated in the introduction and present some further results of [2]. In Section 5.3 we introduce k-profiles of matroids, show how they relate to the matroid's tangles of order k and prove our main result for matroids.

5.1 The k-profiles of graphs

As we have seen in Example 4.3 in Section 4.4, every graph G gives rise to a connectivity system $\mathcal{C}_G$. Let us define a *k-profile of G* as a k-profile of $\mathcal{C}_G$. We first want to show that this notion of a k-profile of G indeed generalizes the notion of a k-block[1]:

Lemma 5.1. *Every $(k-1)$-block b induces a k-profile $P_k(b)$.*

Proof. Let k be a positive integer and let b be a $(k-1)$-block. As in Section 4.1 we obtain a profile from b by

$$P_k(b) := \{(A, B) \in \mathcal{C}_G \mid b \subseteq B \wedge \lambda(A, B) < k\}. \tag{5.1}$$

It is straightforward to check that $P_k(b)$ satisfies (P1) and (P2) and thus is a profile. From the definition it is clear that it is a k-profile. □

In fact, $P_k(b)$ is precisely the $\mathcal{S}_{<k}$-profile of b as a large $\mathcal{S}_{<k}$-block. In what follows we also make use of the language introduced in Section 4.1, in particular we say that b and $P_k(b)$ correspond.

We also promised that the notion of a profile in addition encompasses the notion of a tangle. A *tangle of order k* of G was defined by Robertson and Seymour [14] as a set θ of separations of order less than k of G such that

[1] more precisely: the notion of a $(k-1)$-block.

(θ1) for every separation (A, B) of order less than k of G either (A, B) or (B, A) is contained in θ;

(θ2) for all $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \theta$ we have $G[A_1] \cup G[A_2] \cup G[A_3] \neq G$.

The next lemma shows that tangles of order k are indeed k-profiles.

Lemma 5.2. *Every tangle of order k is a k-profile.*

Proof. Let θ be a tangle of order k. To check (P1) let $(A, B) \in \theta$ and consider $(C, D) \leq (A, B)$. Then we have $G[A] \cup G[D] = G$, which implies $(D, C) \notin \theta$ due to (θ1). Now suppose that θ does not satisfy (P2). Then there are $(A, B), (C, D) \in \theta$ such that $(B \cap D, A \cup C) \in \theta$. But since we have $G[A] \cup G[C] \cup G[B \cap D] = G$ this contradicts (θ1). Hence θ is a profile and due to (θ1) it is a k-profile. $\square$

But not every k-profile is a tangle of order k. For example the k-profile $P_k(b)$ of a $(k-1)$-block b need not be a tangle of order k, as we have seen by Example 3.4. If b is large enough, however, then this is the case, as the following Theorem shows (compare [3, Theorem 6.1]).

Theorem 5.3. *Every $(k-1)$-block b with $|b| > \frac{3}{2}(k-1)$ defines a tangle of order k.*

Proof. Let b be a $(k-1)$-block of more than $\frac{3}{2}(k-1)$ vertices, and let $\theta := P_k(b)$ as defined in (5.1), which is a k-profile due to Lemma 5.1. We show that θ is a tangle of order k. Since θ is a k-profile, it satisfies (θ1). For a proof of (θ2), it suffices to consider three arbitrary separations $(A_1, B_1), (A_2, B_2), (A_3, B_3)$ in θ and show that

$$E(A_1) \cup E(A_2) \cup E(A_3) \not\supseteq E, \tag{$*$}$$

where $E(A_i)$ denotes the set of edges that A_i spans in G.

As $|b| > \frac{3}{2}(k-1)$, there is a vertex $v \in b$ that lies in at most one of the three sets $A_i \cap B_i$, say neither in $A_2 \cap B_2$ nor in $A_3 \cap B_3$. Let us choose v in A_1 if possible. Then, as $b \subseteq B_1$, there is another vertex $w \neq v$ in $b \setminus A_1$. As v and w lie in b, the set $(A_1 \cap B_1) \setminus \{v\}$ does not separate them. Hence there is an edge vu with $u \in B_1 \setminus A_1$. Since $v \notin A_2 \cup A_3$, the edge vu is neither in $E(A_2)$ nor in $E(A_3)$. But vu is not in $E(A_1)$ either, as $u \in B_1 \setminus A_1$, completing the proof of $(*)$. $\square$

Let us show that our notion of K-robustness for k-blocks, given in Chapter 3, is compatible with the notion of K-robustness for k-profiles, given in Chapter 4:

Lemma 5.4. *Let k and K be integers. A $(k-1)$-block b is $(K-1)$-robust as a $(k-1)$-inseparable set, if and only if $P_k(b)$ is K-robust as a k-profile.*

Proof. Let b be a $(k-1)$-block and let $P := P_k(b)$. By the definition of K-robustness and Lemma 3.5, b is $(K-1)$-robust as a $(k-1)$-inseparable set if and only if we have:

> *For every separation (C, D) of order less than k with $b \subseteq D$ and every separation (A, B) of order less than K such that $(A, B) \nparallel (C, D)$ and $|\partial(A \cap D)| < k - 1 > |\partial(B \cap D)|$, we have $b \subseteq A$ or $b \subseteq B$.*

This is easily seen to be equivalent to:

> *For every separation $(C, D) \in P$ and every separation (A, B) of order less than K such that $(A, B) \nparallel (C, D)$ and $\lambda(A \cup C, B \cap D) < ord(P) - 1 > \lambda(B \cup C, A \cap D)$, we have either $(A \cup C, B \cap D) \in P$ or $(B \cup C, A \cap D) \in P$.*

Which in turn is equivalent to $P = P_k(b)$ being K-robust as a k-profile. $\square$

Recall that a k-block is robust if it is K-robust for every $K \in \mathbb{N}$. In that spirit a k-profile is *robust* if it is K-robust for every $K \in \mathbb{N}$. As an immediate corollary of Lemma 5.4 we obtain:

Corollary 5.5. *A $(k-1)$-block is robust if and only if its k-profile is robust.* $\square$

Notice that those $(k-1)$-blocks which are large enough so that their profiles are by Theorem 5.3 guarantied to be tangles, are also large in the sense of Section 3.3. Hence, by Lemma 3.4, the tangles obtained from those large $(k-1)$-blocks are robust. In fact, *all* tangles are robust, as the following lemma shows.

Lemma 5.6. *Every tangle of order k is robust.*

Proof. Let θ be a tangle of order k. Consider a separation $(C, D) \in \theta$ and another separation (A, B) of arbitrary large order such that $(A, B) \nparallel (C, D)$ and $\lambda(A \cup C, B \cap D) < k - 1 > \lambda(B \cup C, A \cap D)$. Suppose that neither $(A \cup C, B \cap D)$ nor $(B \cup C, A \cap D)$ is contained in θ. Then due to (θ1) we have $(B \cap D, A \cup C), (A \cap D, B \cup C) \in \theta$. But then $G[C] \cup G[B \cap D] \cup G[A \cap D] = G$ contradicts (θ2). Hence θ is robust. $\square$

As we have seen, the notion of a k-profile generalizes both the notion of a $(k-1)$-block and the notion of a tangle of order k. Now we want to show that the k-profiles of a graph can be characterised by means of another 'notion of highly-connected piece' of a graph, which was defined by Seymour and Thomas in [15]: Given a set X of vertices of G, a component of $G - X$ is called an *X-flap*; a *haven of order k*, or simply a *k-haven*, is a function h that maps each set X of fewer than k vertices to an X-flap such that $h(X)$ and $h(Y)$ touch each other, for all sets X, Y of $< k$ vertices.

Lemma 5.7. *Every k-profile induces a k-haven.*

Proof. Let P be a k-profile and let X be a set of fewer than k vertices. We shall often make use of the fact that X itself is also 'small' with respect to P:

$$\text{We have } (V, X) \notin P \text{ and therefore } (X, V) \in P. \tag{5.2}$$

Indeed, since $X \subseteq V$ we have $(X, V) \leq (V, X)$ such that (P1) yields $(V, X) \notin P$, and by (4.9) this implies $(X, V) \in P$. Furthermore we shall show:

$$\text{We have } (G - F, X \cup F) \in P \text{ for precisely one } X\text{-flap } F. \tag{5.3}$$

If there exists only one X-flap F, than we have $(G - F, X \cup F) = (X, V) \in P$ due to (5.2). So assume that $\{F_1, \ldots, F_n\}$, with $n \geq 2$, is the set of X-flaps of G. Suppose that we have $(G - F_i, X \cup F_i) \in P$ for more than one $i \in \{1, \ldots, n\}$, say for $i = 1, 2$. Then (P2) yields

$$\big((X \cup F_1) \cap (X \cup F_2), (G - F_1) \cup (G - F_2)\big) = (X, V) \notin P$$

which contradicts (5.2). Now suppose that we have $(G - F_i, X \cup F_i) \notin P$ for all $i \in \{1, \dots, n\}$. Then by (4.9) we $(X \cup F_i, G - F_i) \in P$ for all i. Hence (P2) and (4.9) yield

$$\big((X \cup F_1) \cup (X \cup F_2), (G - F_1) \cap (G - F_2)\big) \in P$$

and by simple induction on n we obtain

$$\big(X \cup \bigcup_{1 \le i \le n} F_i \;,\; G - \bigcup_{1 \le i \le n} F_i\big) = (V, X) \in P,$$

which contradicts (5.2). Hence (5.3) holds.

Now let us define h as the function that maps every set X of fewer than k vertices to the unique X-flap F given by (5.3). In order to show that h is a k-haven it lasts to check

For every two sets X, Y of fewer than k vertices, $h(X)$ and $h(Y)$ touch each other. (5.4)

Suppose that this is false. Then there are two sets X, Y of fewer than k vertices such that $h(X)$ and $h(Y)$ do not touch. Let $X' := \partial h(X)$ denote the boundary of $h(X)$. Then $X' \subseteq X$, so that X' is a set of fewer than k vertices, and $h(X')$ and $h(X)$ are either equal or disjoint. If they are disjoint then (P2) yields

$$\big((X \cup h(X)) \cap (X' \cup h(X')), (G - h(X)) \cup (G - h(X'))\big) = (X'', V) \notin P,$$

with $X'' := X \cap (X' \cup h(X')) \subseteq X$, in contradiction to (5.2). Hence we have $h(X') = h(X)$. Thus we may assume $X = \partial h(X)$ and $Y = \partial h(Y)$. Since $h(X)$ and $h(Y)$ are assumed not to touch, we then have $X \cap h(Y) = \emptyset$, $Y \cap h(X) = \emptyset$ and $h(X) \cap h(Y) = \emptyset$. So again by (P2) we obtain

$$\big((X \cup h(X)) \cap (Y \cup h(Y)), (G - h(X)) \cup (G - h(Y))\big) = (X \cap Y, V) \notin P,$$

which again contradicts (5.2). This shows (5.4) which finishes the proof. □

On the other hand, every k-haven h induces an $\mathcal{S}_{<k}$-orientation $P(h)$: for every separation (A, B) of order $< k$ of G put (A, B) into $P(h)$ if and only if $h(A \cap B)$ is contained in B. It is easy to see that $P(h)$ satisfies (P1), such that $P(h)$ is in fact a $\mathcal{S}_{<k}$-preference. But the converse of Lemma 5.7 need not be true. There may exist a k-haven h in G such that $P(h)$ does not satisfy (P2). With an additional condition on h, however, we can make sure that $P(h)$ is a profile. Let us say that a k-haven h is *tangible* if for each pair of separations $(A, B), (C, D)$ of order less than k such that $h(A \cap B) \subseteq B$, $h(C \cap D) \subseteq D$ and the size of $X := (B \cap D) \cap (A \cup C)$ is less than k we have $h(X) \subseteq B \cap D$.

Then we have the following correspondence between k-profiles and tangible k-havens.

Theorem 5.8. *Every k-profile P induces a tangible k-haven h_P and every tangible k-haven h induces a k-profile $P(h)$, and we have $P(h_P) = P$ and $h_{P(h)} = h$.*

Proof. Let P be a k-profile. By Lemma 5.7, P induces a k-haven h_P, which due to (P2) and (4.9) is tangible. Now let h be a tangible k-haven. We shall show that

$$P(h) := \{(A, B) \mid \lambda(A, B) < k \wedge h(A \cap B) \subseteq B\}$$

is a k-profile. By definition of $P(h)$ we just have to check (P1) and (P2). Suppose that $P(h)$ violates (P1). Then there is $(A,B) \in P(h)$ and $(C,D) \le (A,B)$ with $(D,C) \in P(h)$. But this implies $h(A \cap B) \subseteq B \setminus A$ and $h(C \cap D) \subseteq C \setminus D$ which means that $h(A \cap B)$ and $h(C \cap D)$ cannot touch. The property (P2) follows immediately from the fact that h is tangible.

The stated identities $P(h_P) = P$ and $h_{P(h)} = h$ follow from the following:

A separation (A,B) of order less than k is contained in P if and only if $h_P(A \cap B) \subseteq B$. (5.5)

Consider a separation (A,B) of order less than k, let $X := A \cap B$ and let $F := h_P(X)$. By definition of h_P we have $(G - F, X \cup F) \in P$. Suppose we have $(A,B) \in P$ and $F \subseteq A$. Then we have $(X \cup F, G - F) \le (A,B)$ which by (P1) implies $(G - F, X \cup F) \notin P$, a contradiction. On the other hand, if F is contained in B, we have $(A,B) \le (G - F, X \cup F)$ which implies $(B,A) \notin P$ and therefore $(A,B) \in P$ due to (4.9). This shows (5.5). □

So for every separation (A,B) of order less than k a k-profile P will 'orient' (A,B) towards an $(A \cap B)$-flap. More general, let $\mathcal{S}$ be a separation system and let P be a profile that orients $\mathcal{S}$. We say that P *orients* $\mathcal{S}$ *towards* a set $X \subseteq V$ if for every $(A,B) \in P \cap \mathcal{S}$ we have $X \subseteq B$. As a tangle obtained from a $k \times k$-grid shows, a k-profile need not orient the set of all $(< k)$-separations towards a nonempty set. However, if we consider a nested system of $(< k)$-separations, then this is the case.

Lemma 5.9. *Let $\mathcal{N}$ be a nested system of separations of order $< k$.*

(i) *Every k-profile P orients $\mathcal{N}$ towards a unique large $\mathcal{N}$-block $X(\mathcal{N}, P)$. These are such that $X(\mathcal{N}', P) \supseteq X(\mathcal{N}, P)$ whenever $\mathcal{N}' \subseteq \mathcal{N}$.*

(ii) *$\mathcal{N}$ distinguishes two k-profiles if and only if they orient it towards different large $\mathcal{N}$-blocks.*

Proof. To show (i) consider a k-profile P and let $P_X := P \cap \mathcal{N}$. We want to show that

$$X := \bigcap_{(A,B) \in P_X} B$$

is a large $\mathcal{N}$-block. First we check that X is large. Let (A,B) be a maximal separation in $(P_X, \le)$. Then we have $A \cap B \subseteq X$. Suppose that X is small, that is, $X = (A \cap B)$ and $|X| = k$. Then for every element $b \in B \setminus A$ there is a $\le$-maximal separation $(C_b, D_b) \in P_X$ with $b \in C_b \setminus D_b$ and $C_b \cap D_b = A \cap B$, since otherwise we had $(A \cap B) \subsetneq X$ such that X was large. Applying (P2) to (A,B) and (C_b, D_b) yields $\big(B \cap D_b, A \cup C_b\big) \notin P$. Since $(B \cap D_b) \cap (A \cup C_b) \subseteq (A \cap B) \cup (C_b \cap D_b) = X$ the order of $(B \cap D_b, A \cup C_b)$ is at most k. Thus (4.9) gives $(A \cup C_b, B \cap D_b) \in P$. Using the same argument successively, at each step replacing (A,B) with the separation obtained by the previous step and (C_b, D_b) with an $(C_{b'}, D_{b'})$ for a $b' \in B \setminus A$ that was not considered before such that, in the end, all elements of $B \setminus A$ have been considered, we obtain:

$$\Big(A \cup \bigcup_{b \in B \setminus A} C_b \;,\; B \cap \bigcap_{b \in B \setminus A} D_b\Big) = \big(V(G), X\big) \in P.$$

As we have $\big(X, V(G)\big) \le \big(V(G), X\big)$, this contradicts (P1). So X is large.

Due to (4.9) we know that X is $\mathcal{N}$-inseparable. Consider any element $v \in V(G) \smallsetminus X$. Since v is not contained in X there is $(A,B) \in P_X \subseteq \mathcal{N}$ with $v \in A \smallsetminus B$. On the other hand we have $X \cap (B \smallsetminus A) \neq \emptyset$ since X is large and contained in B. Hence, (A,B) separates $X \cup \{v\}$, which shows that X is maximal $\mathcal{N}$-inseparable and thus a large $\mathcal{N}$-block. By definition of X we have $P_X = P_{\mathcal{N}}(X)$ which shows that P points to X. Suppose that there is another block $X' \neq X$ with $P_{\mathcal{N}}(X') \subseteq P$. Then there is a separation $(E,F) \in \mathcal{N}$ which distinguishes X and X' such that $(E,F) \in P_{\mathcal{N}}(X)$ and $(F,E) \in P_{\mathcal{N}}(X')$. But this implies $(E,F),(F,E) \in P$, in contradiction to (P1). Thus the forst statement of (i) holds. The second statement of (i) follows easily from the definition of X.

For a proof of (ii) consider two distinct k-profiles P, P'. By (i) both P and P' point to a large $\mathcal{N}$-block X and X', respectively, with $P \cap \mathcal{N} = P_{\mathcal{N}}(X)$ and $P' \cap \mathcal{N} = P_{\mathcal{N}}(X')$. Assume first that $\mathcal{N}$ distinguishes P from P'. Then there is $(A,B) \in \mathcal{N}$ with $(A,B) \in P$ and $(B,A) \in P'$. Hence we have $(A,B) \in P_{\mathcal{N}}(X')$ and $(B,A) \in P_{\mathcal{N}}(X)$ such that $P_{\mathcal{N}}(X) \neq P_{\mathcal{N}}(X')$.

For the converse assume $P_{\mathcal{N}}(X) \neq P_{\mathcal{N}}(X')$. Due to (4.9) we have $|P_{\mathcal{N}}(X)| = |P_{\mathcal{N}}(X')|$ so there is $(A,B) \in P_{\mathcal{N}}(X) \smallsetminus P_{\mathcal{N}}(X')$, which once more by (4.9) implies $(B,A) \in P_{\mathcal{N}}(X')$. Hence, (A,B) witnesses that $\mathcal{N}$ distinguishes P and P'. $\square$

With the help of Lemma 5.9 we obtain the following bound on the numbers of k-profiles in G.

Corollary 5.10. *The number of k-profiles of G is at most $|V(G)|$.*

Proof. By (4.11) each pair of k-profiles is K-robustly K-distinguishable, for $K = k$. Hence Theorem 4.16 yields a nested separation system $\mathcal{N}$ that distinguishes each pair of k-profiles of G. Since there can be at most $|V(G)|$ large $\mathcal{N}$-blocks, the given bound follows from Lemma 5.9. $\square$

Note that this bound is asymptotically sharp: consider a star with m leaves and add another vertex joined to all the vertices of the star – this graph has $m+2$ vertices and m distinct 2-profiles.

As another easy consequnece of Lemma 5.9 we obtain the following.

Lemma 5.11. *Let $\mathcal{N}$ be a nested set of separations, and let X be an $\mathcal{N}$-block. If a separation (A,B) that is nested with $\mathcal{N}$ distinguishes two k-profiles that both orient $\mathcal{N}$ towards X, then X is large and (A,B) separates X.*

Proof. Consider two k-profiles P_1 and P_2 that orient $\mathcal{N}$ towards a large $\mathcal{N}$-block X. Let (A,B) be a separation that distinguishes P_1 from P_2 and that is nested with $\mathcal{N}$. We may assume $(A,B) \in P_1$ and $(B,A) \in P_2$. Let $\mathcal{N}' := \mathcal{N} \cup \{(A,B),(B,A)\}$. Since P_1 and P_2 are k-profiles that orient $\mathcal{N}$, we know by (4.9) that all separations in $\mathcal{N}$ have order $< k$. Then by Lemma 5.9 the profiles P_1 and P_2 orient $\mathcal{N}$ towards distinct large $\mathcal{N}'$-blocks X_1 and X_2, respectively. Since both P_1 and P_2 orient $\mathcal{N}$ towards X, we have $X = X_1 \cup X_2$. Furthermore we have $X_1 \subseteq B$ and $X_2 \subseteq A$ and since both X_1 and X_2 are large, we have $X_1 \cap (B \smallsetminus A) \neq \emptyset$ and $X_2 \cap (A \smallsetminus B) \neq \emptyset$. Hence (A,B) separates X. $\square$

5.2 Tree-decompositions distinguishing the k-profiles of a graph

We now come to prove our main results for graphs, Theorems 1.3 to 1.5 from the introduction. We say that a k-block b and a tangle θ of order $k+1$ are *distinguishable* if we have $P_{k+1}(b) \neq \theta$, two k-blocks are *distinguishable* if their $(k+1)$-profiles are, and two tangles are *distinguishable* if they are distinguishable as profiles.

Proof of Theorem 1.3. Let $\mathcal{P}$ be the set of all tangles of order $k+1$ and all k-block profiles. Then for $K := k+1$ every pair of distinguishable k-blocks or tangles of order $k+1$ corresponds to a K-robustly K-distinguishable pair of profiles in $\mathcal{P}$, due to (4.11). Then the theorem follows from Theorem 4.20 for the given integer K, the set of profiles $\mathcal{P}$ and for $\Gamma = \mathrm{Aut}(G)$. $\square$

A tangle θ of order k is *maximal* if for every tangle θ' of order $k' > k$ we have $\theta \not\subseteq \theta'$.

Proof of Theorem 1.4. Let $\mathcal{P}$ be the set of maximal tangles in G and let K be the maximum order of a tangle in $\mathcal{P}$. Then by Lemma 5.6 each pair of profiles in $\mathcal{P}$ is K-robustly K-distinguishable. As in the previous proof, the theorem follows from Theorem 4.20 for the given integer K, the set of profiles $\mathcal{P}$ and for $\Gamma = \mathrm{Aut}(G)$. $\square$

And finally we give a

Proof of Theorem 1.5. Let $\mathcal{P}_k$ be the set of all tangles of order k and robust $(k-1)$-block-profiles of G, let K bet largest integer such that $\mathcal{P}_K \neq \emptyset$ and put $\mathcal{P} := \bigcup_{1 \leq k \leq K} \mathcal{P}_k$. Let $(\mathcal{T}, \mathcal{V})$ be the tree-decomposition obtained from $\mathcal{N}(\mathcal{P})$ as given by Theorem 4.16. Then $(\mathcal{T}, \mathcal{V})$ efficiently distinguishes each pair of distinguishable robust k-blocks and tangles of order $k+1$, for all k simultaneously. $(\mathcal{T}, \mathcal{V})$ is canonical due to Corollaries 2.9 and 4.17. $\square$

For the remainder of the section we concentrate on tree-decompositions which, for fixed $k \in \mathbb{N}$, distinguish all the k-blocks of a graph. In this context, let us call a part of such a tree-decomposition *essential* if it contains a k-block, and *inessential* otherwise. It is clear from the definition that the hubs of our canonical tree-decompositions are inessential parts, but there can be various other kinds of inessential parts. Even essential parts may well contain vertices that are not contained in any k-block, we call the set of all such vertices in the part the *junk* of the part. As we have seen in Section 4.4 we can bound the number of inessential parts by the number of essential parts.

Unfortunately, we cannot bound the amount of junk in an essential part, unless we put further restrictions on the class of graphs under consideration (see Theorem 5.13 and 5.14). The following example shows that there are graphs for which *any* tree-decomposition of adhesion at most k has essential parts which contain arbitrarily much junk.

Example 5.1. Consider a K_5 and enumerate its vertices as $v_1, \dots, v_5$. Now add a K_2 consisting of the vertices x and y, say, connect x with v_1, v_2, v_3 and connect y with v_3, v_4, v_5. The resulting graph, as shown in Figure 5.1, consists of a single 4-block $b = \{v_1, \dots, v_5\}$: the vertex x can be separated from b by

$S_y = \{v_1, v_2, v_3, y\}$ and y can be separated from b by $S_x = \{v_3, v_4, v_5, x\}$. But for every tree-decomposition of adhesion at most 4, the part containing b will contain either x or y as well, since S_x and S_y (more precisely the 'corresponding' separations) are crossing and the only 4-separators of that graph.

We can attach arbitrary many disjoint copies of K_2 in exatly the same way as above, without making the unique 4-block b any larger, but adding arbitrarily much junk to the part containing b.

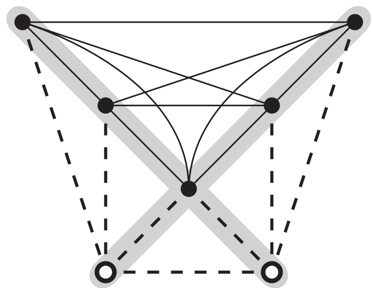

Figure 5.1: The K_5 plus attached junk from Example 5.1.

However, in the case of k-connected graphs, we can entirely isolate those k-blocks that are well-separated, where a k-block b is called *well-separated* if its profile $P_{k+1}(b)$ is well-separated in the the set $\mathcal{S}$ of all proper k-separations.

Let us first prove the following lemma:

Lemma 5.12. *Let G be a k-connected graph, let $\mathcal{S}$ be the set of all its proper k-separations, and let $\mathcal{P}$ be the set of all k-block profiles of G. Then the pair $(\mathcal{S}, \mathcal{P})$ is a task.*

Proof. It is clear that $\mathcal{P}$ is $\mathcal{S}$-distinguishable. Consider two crossing k-separations $(A, B) \nparallel (C, D)$ and two k-blocks b_1 and b_2 such that $(A, B), (C, D) \in P_{k+1}(b_1)$ and $(B, A), (D, C) \in P_{k+1}(b_2)$. Then we have $b_1 \subseteq B \cap D$ and $b_2 \subseteq A \cap C$, which implies that the corresponding corner-separations are proper and therefore both of order at least k (since G is k-connected). So Lemma 3.1 yields that they have order precisely k and are therefore contained in $\mathcal{S}$. $\square$

Theorem 5.13. *Every k-connected graph G has a canonical tree-decomposition of adhesion k such that:*

(i) distinct k-blocks lie in different parts;

(ii) each part containing a well-separated k-block is a k-block;

(iii) if every proper k-separation of G distinguishes two k-blocks, then the part of every leaf is a k-block.

Proof. Let $\mathcal{S}$ be the set of all proper k-separations and let $\mathcal{P}$ be the set of all k-block profiles. Due to Lemma 5.12 the pair $(\mathcal{S}, \mathcal{P})$ is a task. Let σ be any strategy with $\sigma(1) \in \{\mathsf{loc}, \mathsf{max}\}$ and let $(\mathcal{T}, \mathcal{V})$ be the tree-decomposition obtained from $\mathcal{N}_\sigma(\mathcal{S}, \mathcal{P})$. Then Theorem 4.14 and Theorem 2.17 yield (i).

Suppose for a contradiction that there is a part of $(\mathcal{T}, \mathcal{V})$ that contains a well-separated k-block b and a vertex v outside b. Then there exists $(A, B) \in \mathcal{S}$

with $b \subseteq B$ and $v \in A \smallsetminus B$. Among all those separations choose (A, B) such that it is $\leq$-maximal. Then it is also $\leq$-maximal in $P_{k+1}(b)$. Since b is assumed to be well-separated, (A, B) is locally-extremal and therefore contained in $\mathsf{loc}(\mathcal{S}, \mathcal{P})$. Due to (4.6) we have $\mathsf{loc}(\mathcal{S}, \mathcal{P}) \subseteq \mathsf{max}(\mathcal{S}, \mathcal{P})$, such that $(A, B) \in \mathcal{N}_\sigma(\mathcal{S}, \mathcal{P})$. So by Theorem 2.17 the vertices b and v cannot both lie in the same part of $(\mathcal{T}, \mathcal{V})$. This contradiction shows *(ii)*.

To show *(iii)* assume that every proper k-separation of G distinguishes two k-blocks. Then $(\mathcal{S}, \mathcal{P})$ is a reduced task and (4.7) yields:

$$\mathsf{ext}(\mathcal{S}, \mathcal{P}) \subseteq \mathcal{N}_\sigma(\mathcal{S}, \mathcal{P}). \tag{5.6}$$

Now let t be a leaf of $\mathcal{T}$ and let e be the edge of $\mathcal{T}$ incident with t. Let (A, B) be the separation induced by e such that $V_t = B$.

We shall show that (A, B) is extremal in $(\mathcal{S}, \mathcal{P})$. By Lemma 4.5 every $\leq$-maximal element of $\mathcal{S}$ is extremal in $(\mathcal{S}, \mathcal{P})$. So assume for a contradiction that (A, B) is not $\leq$-maximal in $\mathcal{S}$. Then there is an extremal separation $(E, F) \in \mathcal{S}$ with $(A, B) < (E, F)$, such that (E, F) separates two vertices of V_t. But due to (5.6), the separation (E, F) corresponds to an edge e' of $\mathcal{T}$, a contradiction. Hence, (A, B) is $\leq$-maximal, and therefore extremal in $(\mathcal{S}, \mathcal{P})$. So consider any proper k-separation (C, D). By definition of $\mathcal{S}$ we have $(C, D) \in \mathcal{S}$, and since (A, B) is extremal in $(\mathcal{S}, \mathcal{P})$, we have either $(C, D) \leq (A, B)$ or $(D, C) \leq (A, B)$. Hence, we have either $B \subseteq D$ or $B \subseteq C$. Thus, B is k-inseparable and, since B $(= V_t)$ contains a k-block by assumption, it is a k-block, which shows *(iii)*.

By Theorem 4.14 the nested separation system $\mathcal{N}_\sigma(\mathcal{S}, \mathcal{P})$ is canonical, so this is also true for the tree-decomposition $(\mathcal{T}, \mathcal{V})$ obtained from it. And $(\mathcal{T}, \mathcal{V})$ has adhesion k due to Theorem 2.17 and the fact that G is k-connected. □

With an additional condition on a k-connected graph we can even ensure that every part containing a k-block is a k-block.

Theorem 5.14. *Let G be a k-connected graph such that for each edge $e = \{x, y\}$ of G one of the following statements holds:*

(i) the vertices x and y have at least $k - 2$ common neighbours;

(ii) the vertices x and y are joined by at least $\lfloor \frac{3}{2}(k - 1) \rfloor$ independent paths aside from xy;

(iii) the edge e lies in a k-block.

Then G has a canonical tree-decomposition of adhesion k such that every part containing a k-block is a k-block. In particular distinct k-blocks are contained in different parts.

Proof. By Theorem 5.13 it suffices to show for every k-connected graph G satisfying the stated condition that each of its k-blocks is well-separated. So consider such a graph G, let $\mathcal{S}$ be the set of all its proper k-separations and let $\mathcal{P}$ be the set of all its k-block profiles.

Let $P \in \mathcal{P}$ and let $(A, B), (C, D)$ be crossing k-separations in $P \cap \mathcal{S}$. If the order of $(A \cup C, B \cap D)$ is at most k, then it is contained in $P \cap \mathcal{S}$ due to (P2) and (4.9), and we are done. So assume for a contradiction that its order is strictly larger than k. Then by Lemma 3.1 the order of the separation $(B \cup D, A \cap C)$ is strictly less than k. Since G is k-connected, $(B \cup D, A \cap C)$ must be improper. As B and D contain a k-block, there are vertices $b \in B \smallsetminus A$ and $d \in D \smallsetminus C$, and we have $A \cap C \subseteq B \cup D$. Since (A, B) and (C, D) are not nested, we cannot have $A \cap C \subseteq B \cap D$. By symmetry we may assume that there is a vertex $x \in (C \cap D) \smallsetminus B$. Assume further that x has no neighbour $y \in (A \cap B) \smallsetminus D$. Since (C, D) is a proper separation, there must be a vertex $c \in C \smallsetminus D$, which is separated from d by $(C \cap D) \smallsetminus \{x\}$, contradicting the k-connectedness of G.

So suppose, again for a contradiction, that there exists such a neighbour $y \in (A \cap B) \smallsetminus D$ of x and let $e := \{x, y\}$.

Let us first assume that e satisfies (i). All common neighbours of x and y are contained in $A \cap C$, as (A, B) and (C, D) are separations. Thus we have $k \leq |A \cap C| < k$, which is impossible.

Now assume that e satisfies (ii) and let $\mathcal{W}$ be a set of at least $\lfloor \frac{3}{2}(k-1) \rfloor$ independent x-y-paths aside from xy. There are two kinds of paths from x to y: those that meet $(A \cap B) \smallsetminus C$ and those that do not. A path of the first type has to meet also $(C \cap D) \smallsetminus \{x\}$, since it contains a subpath from a vertex in $(A \cap B) \smallsetminus C$ to $y \in (A \cap B) \smallsetminus D$ that avoids x. A path avoiding $(A \cap B) \smallsetminus C$ meets $(A \cap C) \smallsetminus \{x, y\}$. Indeed, $(A \cap C) \smallsetminus \{x, y\}$ together with $(A \cap B) \smallsetminus C$ separates x from y in $G - xy$.

Let U be the set of those vertices in $(A \cap C) \smallsetminus \{x, y\}$ that lie on a path in $\mathcal{W}$. We note that $|U| \leq |(A \cap C) \smallsetminus \{x, y\}| \leq k - 3$. Every path in $\mathcal{W}$ has a vertex either in U or in both $(A \cap B) \smallsetminus (U \cup \{y\})$ and $(C \cap D) \smallsetminus (U \cup \{x\})$. As $\mathcal{W}$ is a set of independent paths, we conclude

$$\begin{aligned}
|\mathcal{W}| &\leq |U| + min\{|(A \cap B) \smallsetminus (U \cup \{y\})|, |(C \cap D) \smallsetminus (U \cup \{x\})|\} \\
&\leq |U| + k - 1 - \lceil |U|/2 \rceil \\
&= k - 1 + \lfloor |U|/2 \rfloor \\
&\leq k - 1 + \lfloor (k-3)/2 \rfloor \\
&= \lfloor \frac{3}{2}(k-1) \rfloor - 1.
\end{aligned}$$

But by (ii), there are more such paths, a contradiction.

So finally assume that e satisfies (iii). Let b be some k-block containing x and y. As $x \in A \smallsetminus B$, the block b has to lie in A and analogously, b lies in C. But then $k + 1 \leq |b| \leq |A \cap C| \leq k - 1$, a contradiction, which finishes the proof. □

Note that every edge of a 2-connected graph obviously satisfies condition (i) of Theorem 5.14 (for $k = 2$). In fact the canonical tree-decomposition of a 2-connected graph obtained by any strategy starting with max corresponds to the well-known decomposition of 2-connected graphs by Tutte [17].

5.3 The k-profiles of matroids

Let M be finite matroid on a groundset E and let $\mathcal{C}_M$ be the connectivity system obtained from M as in Example 4.2. A *k-profile of M* is a k-profile of $\mathcal{C}_M$ and a *tree-decomposition of M* is a tree-decomposition of $\mathcal{C}_M$. As for graphs, an important class of k-profiles of a matroid is given by its tangles of order k.

Following Geelen et al. (see [8, 9]) we define a *tangle of order k* of M as a set θ of separations of order less than k of M that has the following properties:

(θ1) for every separation (A, B) of order $< k$ either (A, B) or (B, A) is in θ;

(θ2) if $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \theta$ then $A_1 \cup A_2 \cup A_3 \neq E$;

(θ3) for every $e \in E$ we have $(E - e, e) \notin \theta$.

And as for graphs we have:

Lemma 5.15. *Every tangle of order k of M is a robust k-profile.*

Proof. Let θ be a tangle of order k of M. Consider a separation $(A, B) \in \theta$ and let $(C, D) \leq (A, B)$. Then we have $A \cup D = E$ and hence $(D, C) \notin \theta$ due to (θ2), which shows (P1). Now consider $(A, B), (C, D) \in \theta$. Then $A \cup C \cup (B \cap D) = E$ so that (θ2) implies $(B \cap D, A \cup C) \notin \theta$, which shows (P2). Together with (θ1) this shows that θ is a k-profile. To show that it is robust consider $(C, D) \in \theta$ and a separation (A, B) of arbitrary order such that $\lambda(A \cup C, B \cap D) < k - 1 > \lambda(B \cup C, A \cap D)$. We have $C \cup (A \cap D) \cup (B \cap D) = E$ which due to (θ2) implies that one of $(A \cap D, B \cup C), (B \cap D, A \cup C)$ is not in θ. But then (θ1) yields that one of $(A \cup C, B \cap D), (B \cup C, A \cap D)$ is in θ, which proves θ to be robust. □

As for graphs a tangle θ of order k is *maximal* if for every tangle θ' of order $k' > k$ we have $\theta \not\subseteq \theta'$. We are now able to prove the last main result stated in the Introduction:

Proof of Theorem 1.6. Let $\mathcal{P}$ be the set of maximal tangles of M and let K be the maximum order of a tangle in $\mathcal{P}$. Then by Lemma 5.15 each pair of profiles in $\mathcal{P}$ is K-robustly K-distinguishable. The theorem follows from Theorem 4.20 for the given integer K, the set of profiles $\mathcal{P}$ and for $\Gamma = \mathrm{Aut}(M)$. □

In the definition of a tangle, Geelen et al. included property (θ3) to avoid a large amount of 'trivial' tangles. In our definition of a k-profile we do not forbid such profiles. So we get another natural class of profiles:

Lemma 5.16. *For every element $e \in E$ and every integer $k \in \mathbb{N}$ the set*

$$P_k(e) := \{(A, B) \mid e \in B \wedge \lambda(A, B) < k\}$$

is a robust k-profile.

Proof. It is easy to see that $P_k(e)$ satisfies (θ1) and (θ2) of the tangle axioms. The proof of Lemma 5.15 does not refer to (θ3), so it carries over. □

We call such a k-profile P, one for which there is an $e \in E$ with $P = P_k(e)$, a *principal* k-profile. A profile that is not principal is called *non-principal*. By Lemma 5.16 we know that there are precisely $|E(M)|$ principal k-profiles of M. With the results of Section 4.4 we can also give an upper bound on the number of non-principal k-profiles:

Theorem 5.17. *There are at most $\frac{1}{2}|E(M)|$ non-principal k-profiles of M.*

Proof. Let $\mathcal{P}$ be the set of all k-profiles of M and let $K = k$. By (4.11) each pair of profiles in $\mathcal{P}$ is K-robustly K-distinguishable. Let $(\mathcal{T}, \mathcal{V})$ be the tree-decomposition given by Theorem 4.20, which distinguishes each pair of profiles in $\mathcal{P}$. We shall show the following:

Every non-principal k-profile of M inhabits a hub-node of degree at least 4 of $(\mathcal{T}, \mathcal{V})$. (5.7)

Suppose this is false. Then there is a non-principal k-profile P that inhabits a part V_t of $(\mathcal{T}, \mathcal{V})$ which is not a hub-node of degree at least 4. Assume first that there is an $e \in V_t$. Then the principal k-profile $P_k(e) \neq P$ will also inhabit V_t, in contradiction to Theorem 4.20 (ii). So V_t is a hub-node of degree n with $1 \leq n \leq 3$. Let $(A_1, B_1), \ldots, (A_n, B_n)$ be the separations in P which are induced by the edges of $\mathcal{T}$ that are incident with t. Then we have $\bigcap_{1 \leq i \leq n} B_i = \emptyset$, since V_t is a hub. If $n = 2$ then we have $(A_1, B_1) = (B_2, A_2)$, in contradiction to (P1). If $n = 3$ then $(B_1 \cap B_2, A_1 \cup A_2) = (A_3, B_3)$ yields a contradiction to (P2). Hence (5.7) holds.

As every part V_t such that t is a leaf of $\mathcal{T}$ has to be a non-empty, there are at most $|E(M)|$ leaves of $\mathcal{T}$. Since every tree has at least twice as many leaves as vertices of degree at least 4, the stated bound follows from (5.7). □

Since a tangle of order k is a non-principal k-profile we easily obtain the following corollary:

Corollary 5.18. *There are at most $\frac{1}{2}|E(M)|$ tangles of order k of M.* □

Bibliography

[1] P. Bellenbaum and R. Diestel. Two short proofs concerning tree-decompositions. *Comb., Probab. Comput.*, 11:1–7, 2002.

[2] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs. In preparation.

[3] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Forcing k-blocks in graphs by minimum and average degree conditions. In preparation.

[4] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. Connectivity and tree structure in finite graphs. *Combinatorica* (to appear).

[5] R. Diestel. *Graph Theory*. Springer, 4th edition, 2010.

[6] M.J. Dunwoody. Cutting up graphs. *Combinatorica*, 2:15–23, 1982.

[7] M.J. Dunwoody and B. Krön. Vertex cuts. arXiv:0905.0064, 2009.

[8] J. Geelen, B. Gerards, N. Robertson, and G. Whittle. Obstructions to branch-decomposition of matroids. *J. Combin. Theory (Series B)*, 96:560–570, 2006.

[9] J. Geelen, B. Gerards, and G. Whittle. Tangles, tree-decompositions and grids in matroids. *J. Combin. Theory (Series B)*, 99:657–667, 2009.

[10] F. Hundertmark. Profiles. An algebraic approach to combinatorial connectivity. arXiv:1110.6207, 2011.

[11] W. Mader. Über den Zusammenhang symmetrischer Graphen. *Arch. Math. (Basel)*, 21:331–336, 1970.

[12] W. Mader. Existenz n-fach zusammenhängender Teilgraphen in Graphen mit genügend hoher Kantendichte. *Abh. Math. Sem. Univ. Hamburg*, 37:86–97, 1972.

[13] J.G. Oxley. *Matroid Theory*. Oxford University Press, 1992.

[14] N. Robertson and P.D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)*, 52:153–190, 1991.

[15] P. Seymour and R. Thomas. Graph searching and a min-max theorem for tree-width. *J. Combin. Theory (Series B)*, 58(1):22 – 33, 1993.

[16] R. Thomas. A menger-like property of tree-width; the finite case. *J. Combin. Theory (Series B)*, 48:67–76, 1990.

[17] W. T. Tutte. *Connectivity in graphs.* Oxford University Press, 1966.

[18] M.E. Watkins. Connectivity of transitive graphs. *J. Combin. Theory*, 8:23–29, 1970.

Name: Fabian Hundertmark
Titel: The tree-like connectivity structure of finite graphs and matroids
Jahr der Drucklegung: 2013

Zusammenfassung

In der vorliegenden Dissertation untersuchen wir die Zusammenhangsstruktur endlicher Graphen und Matroide. Wir zeigen, dass sich die unterscheidbaren hochzusammenhängenden Teile eines Graphen oder Matroids als unterschiedliche Orientierungen seiner Teilungen, die wir als *Profile* bezeichnen, beschreiben lassen. Insbesondere können wir die maximalen k-untrennbaren Eckenmengen eines Graphen, seine k-*Blöcke*, durch Profile (der Ordnung $k+1$) beschreiben. Ferner zeigen wir, dass die Tangles der Ordnung k eines Graphen oder Matroids spezielle k-Profile sind. Als Hauptresultat dieser Arbeit zeigen wir, dass zu jedem Graphen oder Matroid und zu jedem $k \in \mathbb{N}$ eine kanonische, d.h. nur von der Struktur des Graphen oder Matroids abhängende, Baumzerlegung der Adhäsion kleiner k existiert, die alle k-Profile des Graphen oder Matroids unterscheidet. Anschaulich bedeutet dies, dass die hochzusammenhängenden Teile eines Graphen oder Matroids untereinander baumartig verbunden sind, also dass jeder Graph und jedes Matroid eine baumartige Zusammenhangsstruktur aufweist.

In Kapitel 2 zeigen wir zunächst, dass die Baumzerlegungen einer beliebigen endlichen Menge durch verschachtelte Systeme von Teilungen beschrieben werden können, und umgekehrt, dass jedes verschachtelte System von Teilungen einer Menge durch einen Baum, bzw. eine Baumzerlegung dieser Menge, beschrieben wird. Außerdem betrachten wir, unter welchen Bedingungen aus einem nicht verschachtelten System von Teilungen ein verschachteltes Teilsystem mit ähnlichen Trennungseigenschaften ausgewählt werden kann.

In Kapitel 3 beschäftigen wir uns mit den k-Blöcken eines Graphen. Wir zeigen, dass diese durch eine Baumzerlegung kleiner Adhäsion unterschieden werden können, und untersuchen, wie die Existenz von k-Blöcken in einem Graphen mit anderen Invarianten des Graphen zusammenhängt.

Im vierten Kapitel führen wir Profile ein und diskutieren, unter welchen Bedingungen eine Menge von Profilen durch ein verschachteltes System von Teilungen unterschieden werden kann. Wir diskutieren, wie der Zusammenhang eines Graphen oder Matroids durch eine Bewertung geeigneter Teilungen seiner Grundmenge beschrieben werden kann. In Kapitel 5 wenden wir schließlich die allgemeinen Resultate aus Kapitel 4 auf Graphen und Matroide an.

Lebenslauf

Fabian Hundertmark, geboren am 14.04.1983 in Hamburg

Wissenschaftliche Ausbildung

10/2009 – 09/2012	Promotionsstudium an der Universität Hamburg
04/2009	Diplom in Mathematik
10/2003 – 04/2009	Studium Mathematik an der Universität Hamburg
06/2002	Abitur
08/1993 – 06/2002	Gymnasium Oldenfelde, Hamburg
08/1989 – 06/1993	Grundschule Islandstraße, Hamburg

Stipendien und Anstellungen

seit 10/2011	Wissenschaftlicher Mitarbeiter an der Universität Hamburg
10/2009 – 09/2011	Promotionsstipendium der Universität Hamburg
10/2008 – 03/2009	Studentische Hilfskraft, Universität Hamburg
10/2007 – 03/2008	Studentische Hilfskraft, Universität Hamburg
10/2005 – 07/2006	Studentische Hilfskraft, Universität Hamburg
09/2002 – 06/2003	Zivildienst, Evangelische Jugendhilfe Hamburg