University of Hamburg
Faculty of Mathematics, Informatics and Natural Sciences
Fachbereich Mathematik

Master of Science Thesis

Polishing tree-decompositions to bring out the k-blocks

by

Jochen Pascal Gollin

Supervisor: Prof. Dr. Reinhard Diestel

Hamburg, 2014

Die vorliegende Arbeit habe ich selbständig verfasst und keine anderen als die angegebenen Hilfsmittel - insbesondere keine im Quellenverzeichnis nicht benannten Internet-Quellen - benutzt.

Die Arbeit habe ich vorher nicht in einem anderen Prüfungsverfahren eingereicht.

Die eingereichte schriftliche Fassung entspricht genau der auf dem elektronischen Speichermedium.

Polishing tree-decompositions to bring out the k-blocks

J. Pascal Gollin

October 21, 2014

Abstract

A k-*block* in a graph G is a maximal set of at least k vertices no two of which can be separated in G by removing less than k vertices. It is *polishable* if it appears as a part of some tree-decomposition of adhesion less than k of G.

Extending results of Carmesin, Diestel, Hamann, Hundertmark and Stein, we construct for any finite graph a canonical tree-decomposition of adhesion less than k distinguishing the k-blocks and the tangles of order k with the additional property that every polishable k-block is equal to the unique part in which it is contained. This proves a conjecture of Diestel.

Contents

1 Introduction **2**

2 Preliminaries **3**

3 Construction methods **10**

3.1 Sticking tree-decompositions together 10

3.2 Obtaining tree-decompositions from almost nested sets of separations 13

3.3 Extending a nested set of separations 16

4 Proof of the main result **19**

5 Comparison with other tree-decompositions **22**

6 Appendix **23**

1 Introduction

Robertson and Seymour [10] proved that every finite graph has a tree-decomposition distinguishing any two maximal tangles. Carmesin, Diestel, Hamann and Hundertmark [3] constructed such tree-decompositions in a canonical way that distinguish the k-profiles[1] of the graph for any $k \in \mathbb{N}$.[2] Like tangles, the k-profiles can be thought of as "highly connected pieces" of a graph. While every tangle of order k is a k-profile, there are k-profiles that are not tangles. A *k-block* in a graph G is a maximal set of at least k vertices no two of which can be separated in G by removing less than k vertices. Every k-block induces a k-profile as well.[3]

In [3], Carmesin et al. constructed several tree-decompositions of a finite graph each of which has adhesion less than k and distinguish any two k-profiles. They are also canonical in that they are invariant under the automorphisms of G. Carmesin et al. [4] gave examples of graphs where the size of some parts of any canonical tree-decomposition needs to be much larger than the tree-width of the graph. For applications of tree-decompositions in general it has turned out to be useful that its parts are no larger than some aspect of the graph forces them to be. Thomas [12] introduced the notion of a lean tree-decomposition and showed that every graph G has a lean tree-decomposition witnessing the tree-width of G. These tree-decompositions need not be canonical and there are graphs without any canonical lean tree-decomposition, as [4, Example 1] demonstrates.

In this paper we want to recover some aspect of leanness for canonical tree-decompositions of adhesion less than k that distinguish any two k-profiles: we want some parts of the tree-decomposition to contain no unnecessary vertices. For parts containing k-blocks this notion is easy to capture: Let us call a k-block *polishable* if it is equal to a part of some tree-decomposition of adhesion less than k of G.

The following theorem is the main result of this paper and was conjectured by Diestel [7] (see also [4]).

Theorem 1.1. *Every finite graph has a canonical tree-decomposition $\mathcal{T}$ of adhesion less than k distinguishing any two k-profiles such that every polishable k-block is equal to the unique part in which it is contained.*

Our main result implies a structural characterization of polishability.

[1] The precise definition of a k-profile is given in Section 2.

[2] This extended earlier results of Carmesin, Diestel, Hundertmark and Stein [5].

[3] For more details on k-profiles, tangles and k-blocks and on how these notions relate, see [1]. For more details on profiles in general, see [9].

Corollary 1.2. *A k-block b of a finite graph G is polishable, if and only if the neighbourhood of each component of $G - b$ in G has size less than k.*

We also show that every polishable k-block appears as a part in any lean tree-decomposition of adhesion less than k. Hence our main construction captures that aspect of leanness for canonical tree-decompositions of adhesion less than k.

After recalling some preliminaries in Section 2, we develop the necessary tools in Section 3. Section 4 then uses these tools to prove Theorem 1.1. In Section 5 we will compare the tree-decompositions we constructed with other tree-decompositions. Finally, in Section 6 we give an alternate characterization of a notion from Section 3.

The results of this dissertation are published in [6].

2 Preliminaries

Unless otherwise mentioned, G will always denote a finite, simple and undirected graph with vertex set $V(G)$ and edge set $E(G)$. Any graph-theoretic term and notation not defined here are explained in [8]. First we recall some notations from [3].

An ordered pair (A, B) of subsets of $V(G)$ is a *separation* of G if $A \cup B = V(G)$ and if there is no edge $e = vw \in E(G)$ with $v \in A \setminus B$ and $w \in B \setminus A$. The cardinality $|A \cap B|$ of the *separator* $A \cap B$ of a separation (A, B) is the *order* of (A, B) and a separation of order k is called a *k-separation*. A separation (A, B) is *proper* if neither $A \subseteq B$ nor $B \subseteq A$. Otherwise (A, B) is *improper*. A separation (A, B) is *tight* if every vertex in $A \cap B$ has a neighbour in $A \setminus B$ and a neighbour in $B \setminus A$.

The set of separations of G is partially ordered via

$$(A, B) \leq (C, D) \;:\Leftrightarrow\; A \subseteq C \;\wedge\; D \subseteq B.$$

For no two proper separations (A, B) and (C, D), the separation (A, B) is $\leq$-comparable with (C, D) and (D, C). In particular we obtain that (A, B) and (B, A) are not $\leq$-comparable.

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

$$(A, B) \leq (C, D) \;\Leftrightarrow\; (D, C) \leq (B, A),$$

being nested is symmetric and reflexive. Separations that are not nested are called *crossing*.

A separation (A, B) is *nested* with a set S of separations if (A, B) is nested with every $(C, D) \in S$. A set S of separations is *nested* with a set S' of separations if every $(A, B) \in S$ is nested with S' or equivalently every $(C, D) \in S'$ is nested with S.

A set N of separations is *nested* if any two of its elements are nested. A set S of separations is *symmetric* if for every $(A, B) \in S$ it also contains its *flip* separation (B, A). A symmetric set S of separations is also called a *separation system* or a *system of separations*, and if all its separations are proper, S is called a *proper separation system*. For a set S of separations the separation system *generated by* S is the separation system consisting of the separations in S and their flips. A set S of separations is *canonical* if it is invariant under the automorphisms of G, i.e. for every $(A, B) \in S$ and for every $\varphi \in \mathrm{Aut}(G)$ we obtain $(\varphi[A], \varphi[B]) \in S$.

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

For $k \in \mathbb{N}$ let $S_{<k}$ denote the set of separations of order less than k of G. The *$(< k)$-inseparable* sets are the $S_{<k}$-inseparable sets. So the *k-blocks* are exactly the $S_{<k}$-blocks of size at least k.

For two separations (A, B) and (C, D) not equal to $(V(G), V(G))$ consider a *cross-diagram* as in Figure 1. Every pair $(X, Y) \in \{A, B\} \times \{C, D\}$ denotes a *corner* of this cross-diagram, which we also denote with $\mathrm{cor}(X, Y)$. Let $\overline{X} \in \{A, B\} \setminus \{X\}$ and $\overline{Y} \in \{C, D\} \setminus \{Y\}$. In the diagram we consider the *center* $c \coloneqq A \cap B \cap C \cap D$ and for a corner $\mathrm{cor}(X, Y)$ as above the *interior* $\mathrm{int}(X, Y) \coloneqq (X \cap Y) \setminus (\overline{X} \cup \overline{Y})$ and the *links* $\ell_X \coloneqq (X \cap Y \cap \overline{Y}) \setminus c$ and $\ell_Y \coloneqq (Y \cap X \cap \overline{X}) \setminus c$. The vertex set $X \cap Y$ is the disjoint union of $\mathrm{int}(X, Y)$ with ℓ_X, ℓ_Y and c and thus can be associated with the corner $\mathrm{cor}(X, Y)$.

Remark 2.1. *Two separations (A, B) and (C, D) are nested, if and only if for one of their corners $\mathrm{cor}(X, Y)$ the interior $\mathrm{int}(X, Y)$ and its links ℓ_X and ℓ_Y are empty.* □

For a corner $\mathrm{cor}(X, Y)$ there is a *corner separation* $(X \cap Y, \overline{X} \cup \overline{Y})$, which is again a separation of G.

Lemma 2.2. [5, Lemma 2.2] *For two crossing separations (A, B) and (C, D) any of its corner separation is nested with every separation that is nested with both (A, B) and (C, D).*

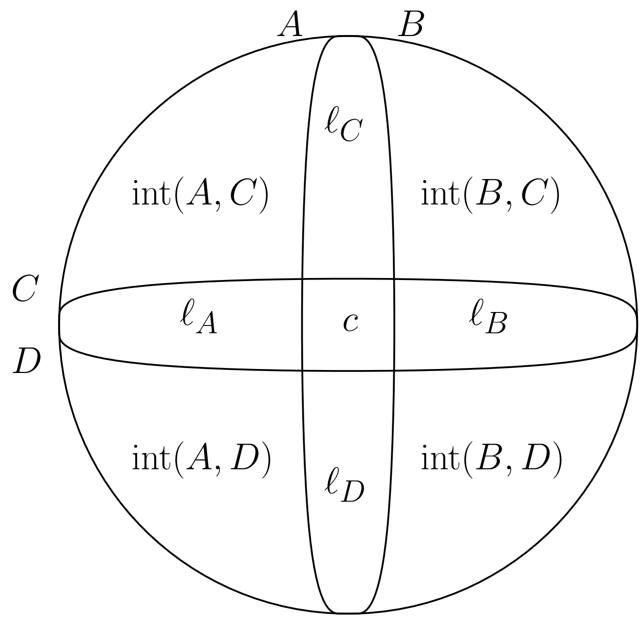

Figure 1: cross-diagram for (A, B) and (C, D)

In particular a corner separation is nested with (A, B), (C, D) and all corner separations.

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

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

(T2) for every edge $e \in E(G)$ there is a node $t \in V(T)$ such that both end vertices of e lie in P_t;

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

The sets P_t are the *parts* of $\mathcal{T}$. For an edge $tt' \in E(T)$ the intersection $P_t \cap P_{t'}$ is the corresponding *adhesion set* and the maximum size of an adhesion set of $\mathcal{T}$ is the *adhesion* of $\mathcal{T}$. A node $t \in V(T)$ is a *hub node* if the corresponding part P_t is a subset of $P_{t'}$ for some neighbour t' of t. If t is a hub node, then P_t is a *hub*. A tree-decomposition $\mathcal{T} = \big(T, (P_t)_{t \in V(T)}\big)$ of G and a tree-decomposition $\mathcal{T}' = \big(T', (P'_t)_{t \in V(T')}\big)$ of G' are *isomorphic* if there is a graph isomorphism $\varphi : G \to G'$ that induces for every part of $\mathcal{T}$ an isomorphism between that part and a part of $\mathcal{T}'$ and induces an isomorphism between T and T'. We say φ *induces* an isomorphism between $\mathcal{T}$ and $\mathcal{T}'$. A tree-decomposition $\mathcal{T}$ is *canonical* if it is invariant under the automorphisms of G, i.e. every automorphism of G induces an automorphism of $\mathcal{T}$.

Let $\big(T, (P_t)_{t\in V(T)}\big)$ be a tree-decomposition of G. For $t \in V(T)$ the *torso* H_t is the graph obtained from $G[P_t]$ by adding all edges joining two vertices in a common adhesion set $P_t \cap P_u$ for any $tu \in E(T)$. A separation (A, B) of $G[P_t]$ is a separation of H_t if and only if it does not separate any adhesion set $P_t \cap P_{t'}$ for $tt' \in E(T)$. A separation (A, B) of G with $A \cap B \subseteq P_t$ for some node $t \in V(T)$ that does not separate any adhesion set $P_t \cap P_{t'}$ for $tt' \in E(T)$ *induces* the separation $(A \cap P_t, B \cap P_t)$ of H_t.

Every oriented edge $\vec{e} = t_1t_2$ of T divides $T - e$ in two components T_1 and T_2 with $t_1 \in V(T_1)$ and $t_2 \in V(T_2)$. By [8, Lemma 12.3.1] e *induces* the separation $\big(\bigcup_{t\in V(T_1)} P_t, \bigcup_{t\in V(T_2)} P_t\big)$ of G such that the separator coincides with the adhesion set $P_{t_1} \cap P_{t_2}$. We say a separation is *induced* by $\mathcal{T}$ if it is induced by an oriented edge of T.

The set of separations induced by a tree-decomposition $\mathcal{T}$ (of adhesion less than k) is a nested system $N(\mathcal{T})$ of separations (of order less than k). We say $N(\mathcal{T})$ is *induced* by $\mathcal{T}$. Clearly if $\mathcal{T}$ is canonical, then so is $N(\mathcal{T})$. Conversely, as proven in [5], every nested separation system N *induces* a tree-decomposition $\mathcal{T}(N)$:

Theorem 2.3. [5, Theorem 4.8] *Let N be a canonical nested separation system of G. Then there is a canonical*[4] *tree-decomposition $\mathcal{T}(N)$ of G such that*

(i) every N-block of G is a part of $\mathcal{T}(N)$;

(ii) every part of $\mathcal{T}(N)$ is either an N-block of G or a hub;

(iii) the separations of G induced by $\mathcal{T}(N)$ are precisely those in N,

(iv) every separation in N is induced by a unique oriented edge of $\mathcal{T}(N)$.

Let S be a separation system. A subset $O \subseteq S$ is an *orientation* of S, if for every $(A, B) \in S$ exactly one of (A, B) and (B, A) is an element of O. An orientation O of S is *consistent*, if for every (A, B), $(C, D) \in S$ with $(A, B) \in O$ and $(C, D) \le (A, B)$ we obtain $(C, D) \in O$ as well.[5]

[4] In the original paper this theorem is stated without the word canonical because it holds in a greater generality. But it is clear from the proof that if N is canonical, then so is $\mathcal{T}$.

[5] In other contexts consistency is defined by requireing $(D, C) \notin O$, which is in our context equivalent.

A consistent orientation Q of $S_{<k}$ is called a *k-profile* or a *profile of order k* if it satisfies

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

In particular if the order $|(A \cup C) \cap (B \cap D)|$ of this corner separation is less than k, we have $(A \cup C, B \cap D) \in Q$.

It is easy to check that every k-block b *induces* a k-profile via

$$P_k(b) := \big\{(A,B) \in S_{<k} \;\big|\; b \subseteq B\big\}.$$

Also *k-tangles*, as introduced by Robertson and Seymour [10], are k-profiles. For more background on profiles, see [1] or [9].

For $r \in \mathbb{N}$, a k-profile Q is *r-robust* if for any $(A,B) \in Q$ and any $(C,D) \in S_{<r+1}$ one of $(A \cup C, B \cap D)$, $(A \cup D, B \cap C)$ either has order at least $k-1$, or is in Q.

Lemma 2.4. *(i) Every k-profile is ℓ-robust for all $\ell < k$;*

(ii) if a k-block b contains a complete graph on k vertices, then the induced profile $P_k(b)$ is r-robust for all $r \in \mathbb{N}$.

Proof. (i) is a direct consequence of (P). For (ii), let $(A,B) \in P_k(b)$ and $(C,D) \in S_{<r+1}$. Suppose for a contradiction that the relevant corner separations have order less than $k-1$ and are both not in $P_k(b)$. Hence $(B \cap D, A \cup C)$ and $(B \cap C, A \cup D)$ are in $P_k(b)$ and therefore $b \subseteq (A \cap B) \cup \ell_B$. But since each of the separators of the two corner separations cannot contain the complete subgraph K_k of b, both ℓ_C and ℓ_D contain a vertex of K_k, contradicting that (C,D) is a separation. □

Lemma 2.5. [6] *Let $X \subseteq V(G)$ with $|X| < k$ and let Q be a k-profile. Then there exists a component C of $G - X$ such that $(V(G) \setminus C, C \cup X) \in Q$. Furthermore, $(V(G) \setminus C, C \cup N(C)) \in Q$ as well.*

Proof. Let $C_1, \dots, C_n$ denote the components of $G - X$ and for $i \in \{1, \dots, n\}$ let $(A_i, B_i) := (V(G) \setminus C_i, C_i \cup X)$. To reach a contradiction suppose that $(B_i, A_i) \in Q$ for all $i \in \{1, \dots, n\}$. Then (P) yields inductively for all $m \leq n$ that $\big(\bigcup_{i \leq m} B_i, \bigcap_{i \leq m} A_i\big) \in Q$, since their separators all equal X. Hence for $m = n$, we obtain $(V(G), X) \in Q$, contradicting the consistency of Q with $(X, V(G)) \leq (V(G), X)$. Thus there is a component C of $G - X$ such that $(A,B) := (V(G) \setminus C, C \cup X) \in Q$.

[6] This lemma basically states, that every k-profile induces a k-haven, as introduced by Seymour and Thomas [11]. For more details again see [1].

Now suppose $(C \cup N(C), V(G) \setminus C) \in Q$. Then (P) with (A, B) yields $\big((V(G) \setminus C) \cup C \cup N(C), (C \cup X) \cap (V(G) \setminus C)\big) = (V(G), X) \in Q$, contradicting the consistency of Q again. Hence $(V(G) \setminus C, C \cup N(C)) \in Q$. □

A k-profile Q *inhabits* a part P_t of a tree-decomposition $\big(T, (P_t)_{t \in V(T)}\big)$ if for every $(A, B) \in Q$ we obtain that $(B \setminus A) \cap P_t$ is not empty.

Corollary 2.6. *Let $\big(T, (P_t)_{t \in V(T)}\big)$ be a tree-decomposition and let Q be a k-profile. If Q inhabits a part P_t, then $|P_t| \geq k$.*

Proof. Our aim is to show that if $|P_t| < k$, then any profile Q does not inhabit P_t. By Lemma 2.5 there is a component C of $G - P_t$ such that $(V(G) \setminus C, C \cup P_t) \in Q$. Since $(C \cup P_t) \setminus (V(G) \setminus C) = C$ and since $C \cap P_t$ is empty, we obtain that Q does not inhabit P_t. □

Note that if for a node $t \in V(T)$ every separation induced by an oriented edge ut of T has order less than k, then Q inhabits P_t if and only if all those separations are in Q.

Lemma 2.7. *Let Q be a profile of order at most k and let $\mathcal{T} = \big(T, (P_t)_{t \in V(T)}\big)$ be a tree-decomposition of adhesion less than k. Then there is a part P_t of $\mathcal{T}$ inhabited by Q.*

Proof. If $\mathcal{T}$ is trivial, then clearly Q inhabits $V(G)$. Hence let $\mathcal{T}$ be non-trivial. Let $(C, D) \in N(\mathcal{T})$ be maximal such that there is a separation $(A, B) \in Q$ with $(C, D) \leq (A, B)$. If no such (C, D) exists, let $(C, D) \in N(\mathcal{T})$ be arbitrary. Let $t \in V(T)$ be such that (C, D) is induced by an incoming edge to t and no outging edge from t, which exists since T is finite. We show that Q inhabits P_t.

Suppose for a contradiction that Q does not inhabit P_t. Let $(A, B) \in Q$ witness this, i.e. $(B \setminus A) \cap P_t$ is empty. By Lemma 2.5 there is a component K of $G - (A \cap B)$ such that $(V(G) \setminus K, K \cup N(K)) \in Q$. By consistency we obtain $K \subseteq (B \setminus A)$, and hence $K \cap P_t$ is empty. Thus for all (E, F) induced by any outgoing edge from t we obtain either $K \subseteq (F \setminus E)$ or $K \subseteq (E \setminus F)$, since $(E \cap F) \subseteq P_t$. If $K \subseteq E \setminus F$ for all such (E, F), then by construction of $N(\mathcal{T})$ we obtain $K \subseteq P_t$, contradicting that $K \cap P_t$ is empty. Hence there is a separation (E, F) induced by an outgoing edge from t such that $K \subseteq (F \setminus E)$. Therefore we obtain $(E, F) \leq (V(G) \setminus K, K \cup N(K))$. Hence (E, F) is a separation induced by $\mathcal{T}$ such that there is a separation $(A', B') \in Q$ with $(E, F) \leq (A', B')$. Thus (C, D) was indeed chosen maximal with that property.

If $(E,F) = (D,C)$, then $(K \cup N(K), V(G) \setminus K) \leq (C,D) \leq (A',B')$ for the $(A',B') \in Q$ which is greater or equal than (C,D), contradicting the consistency of Q. And since $(C,D) \neq (E,F)$ by the choice of t, we obtain $(C,D) < (E,F)$, contradicting the maximality of (C,D). □

Two profiles Q and Q' of order at most k are *k-distinguishable* if there is a separation $(A,B) \in S_{<k}$ with $(A,B) \in Q$ and $(B,A) \in Q'$. Such a separation *distinguishes* Q and Q'. It is said to distinguish Q and Q' *efficiently* if the order $|A \cap B|$ is minimal among all separations in $S_{<k}$ which distinguish Q and Q'. A set S of separations of order less than k *distinguishes* a set $\mathcal{P}$ of profiles of order at most k (efficiently) if any two distinct $Q, Q' \in \mathcal{P}$ are distinguished by some $(A,B) \in S$ (efficiently). A tree-decomposition $\mathcal{T}$ *distinguishes* a set $\mathcal{P}$ of profiles of order at most k if every part of $\mathcal{T}$ is inhabited by at most one profile of $\mathcal{P}$. It distinguishes $\mathcal{P}$ *efficiently* if any two distinct $Q, Q' \in \mathcal{P}$ are efficiently distinguished by some (A,B) induced by $\mathcal{T}$. It is easy to verify that a tree-decomposition $\mathcal{T}$ distinguishes a set $\mathcal{P}$ of profiles of order at most k efficiently, if and only if $N(\mathcal{T})$ does.

For our main result of this paper, we will build on the following theorem.

Theorem 2.8. *Every graph G has a tree-decomposition $\mathcal{T}$ of adhesion less than k distinguishing any two k-distinguishable $(k-1)$-robust profiles of order at most k of G efficiently.*

We obtain this version of the theorem by combining a result of [2] together with some methods of Section 3. The notion of a *tree of tree-decompositions* introduced there directly translates into *almost nested* sets of separations developed here. Hence Theorem 2.8 follows from Theorem 3.7 and [2, Theorem 9.2]. We will give more details on this characterization in the Appendix. Theorem 2.8 also follows from a result of Hundermark and Lemanczyk [9]. This extends earlier versions of Carmesin, Diestel, Hundertmark, Hamann and Stein [5, Theorem 6.5] and [3, Theorem 4.4].

A vertex is called *central* in G if the greatest distance to any other vertex is minimal. It is well known that a finite tree T has either a unique central vertex or precisely two central adjacent vertices v and w. In the second case vw is called a *central edge*. For a vertex or edge to be central is obviously a property invariant under automorphisms of G.

3 Construction methods

3.1 Sticking tree-decompositions together

Given a tree-decomposition $\mathcal{T}$ of G and for each torso H_t of $\mathcal{T}$ a tree-decomposition of H_T we want to construct a new tree-decomposition $\overline{\mathcal{T}}$ of G by gluing together the tree-decompositions of the torsos along $\mathcal{T}$ in a canonical way.

Example 3.1. In this example we explain the canonical gluing for a particular graph. Consider the graph G obtained from three disjoint triangles D_1, D_2, D_3 by first identifying an edge of D_2 with an edge of D_3 and then identifying a vertex of that edge with a vertex of D_1. In Figure 2, we depicted G together with one of its tree-decompositions, drawn in black. For the upper torso we take the trivial tree-decomposition whilst for the lower torso we take the tree-decomposition drawn in gray, whose tree is just an edge. In order to stick these tree-decompositions of the torsos together in a canonical way, which must be invariant under the automorphism group of G, we first have to add a hub node to the gray tree-decomposition, see Figure 3 for the stuck together tree-decomposition.

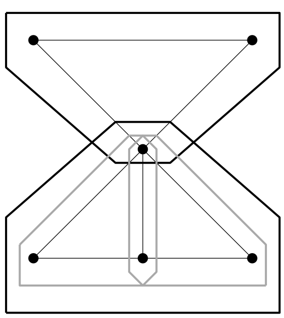

Figure 2: G with a tree-decomposition

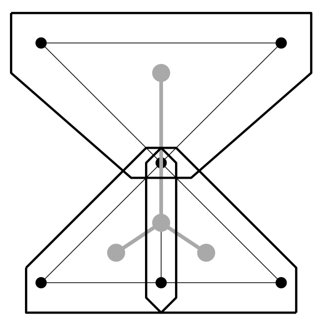

Figure 3: G with canonical gluing

Before we can construct $\overline{\mathcal{T}}$, we need some preparation. Given a tree-decomposition $\mathcal{T} = \big(T, (P_t)_{t\in V(T)}\big)$ we construct a new tree-decomposition $\widetilde{\mathcal{T}} = \big(\widetilde{T}, (\widetilde{P}_t)_{t\in V(\widetilde{T})}\big)$ by contracting every edge tu of T where $P_t = P_u$. In this tree-decomposition two adjacent nodes never have the same part. Let $F \subseteq E(\widetilde{T})$ be the set of edges tu where neither $\widetilde{P}_t \subseteq \widetilde{P}_u$ nor $\widetilde{P}_u \subseteq \widetilde{P}_t$. By subdiving every edge $tu \in F$ and assigning to the subdivided node the part $\widetilde{P}_t \cap \widetilde{P}_u$, we obtain a new tree-decomposition $\widehat{\mathcal{T}} = \big(\widehat{T}, (\widehat{P}_t)_{t\in V(\widehat{T})}\big)$.

Remark 3.2. *$\widehat{\mathcal{T}}$ satisfies the following:*

(i) for every node $t \in V(T)$ there is a node $u \in V(\widehat{T})$ such that $P_t = \widehat{P}_u$;

(ii) every separation induced by $\widehat{\mathcal{T}}$ is also induced by $\mathcal{T}$;

(iii) for every edge $tu \in E(\widehat{T})$ precisely one of $\widehat{P}_t$ or $\widehat{P}_u$ is a proper subset of the other;

(iv) for every edge $tu \in E(T)$ that induces a separation not induced by $\widehat{\mathcal{T}}$ we obtain $P_t = P_u$.

If $\mathcal{T}$ is canonical, then $\widehat{\mathcal{T}}$ is canonical as well. □

Lemma 3.3. *Let $\big(T, (P_t)_{t\in V(T)}\big)$ be a tree-decomposition of a graph G and K be a complete subgraph of G. Then there is a node t of $\widehat{T}$ such that $K \subseteq \widehat{P}_t$ and its definition is invariant under the automorphism group of G.*

Proof. As K is complete, there is a node $u \in V(\widehat{T})$ with $K \subseteq \widehat{P}_u$ by (T2).

In order to define a node t with this property in a canonical way, let W be the subforest of T consisting of those nodes w with $K \subseteq \widehat{P}_w$, which is connected by (T3). Now W either has a central vertex t or a central edge tu such that $\widehat{P}_u$ is a proper subset of $\widehat{P}_t$. In both cases the definition of t is invariant under the automorphism group of G. □

Let $\mathcal{T} = \big(T, (P_t)_{t\in V(T)}\big)$ be a canonical tree-decomposition of G. Two torsos H_t and H_u of $\mathcal{T}$ are *similar*, if there is an automorphism of G that induces an isomorphism between H_t and H_u. A family $\big(\mathcal{T}^t\big)_{t\in V(T)}$ where $\mathcal{T}^t = \big(T^t, (P^t_u)_{u\in V(T_t)}\big)$ is a canonical tree-decomposition of the torso H_t for each $t \in V(T)$ is *canonical* if for any two similar torsos H_t and H_u of $\mathcal{T}$ the automorphism of G that witnesses the similarity of H_t and H_u induces an isomorphism between $\mathcal{T}^t$ and $\mathcal{T}^u$.

Lemma 3.4. *Let $\mathcal{T} = \big(T, (P_t)_{t\in V(T)}\big)$ be a canonical tree-decomposition of G and let $\big(\mathcal{T}^t\big)_{t\in V(T)}$ be a canonical familiy of tree-decompositions, where $\mathcal{T}^t = \big(T^t, (P^t_u)_{u\in V(T_t)}\big)$ is a tree-decomposition of the torso H_t for $t \in V(T)$. Then there is a canonical tree-decomposition $\overline{\mathcal{T}} = \big(\overline{T}, (\overline{P}_t)_{t\in V(\overline{T})}\big)$ of G such that*

(i) for $t \in V(T)$ every node $u \in V(T^t)$ is also a node of $\overline{T}$ and $\overline{P}_u = P^t_u$;

(ii) every node $u \in V(\overline{T})$ that is not a node of any T^t is a hub node;

(iii) *every separation induced by $\overline{\mathcal{T}}$ is either induced by $\mathcal{T}$ or induces a separation of a torso H_t, which is induced by $\mathcal{T}^t$;*

(iv) *every separation induced by $\mathcal{T}$ is also induced by $\overline{\mathcal{T}}$;*

(v) *if two profiles Q_1 and Q_2 of order at most k of a torso H_t are efficiently distinguished by $\mathcal{T}^t$, then there is a separation induced by $\overline{\mathcal{T}}$ that induces a separation on H_t which also distinguishes Q_1 and Q_2 efficiently.*

Proof. For every tree-decomposition $\mathcal{T}^t$ consider $\widehat{\mathcal{T}}^t$ as in Remark 3.2. For $e = tu \in E(T)$, let A_e be the adhesion set $P_t \cap P_u$. Since $H_t[A_e]$ is complete, there is a canonically defined node $\gamma(t,u)$ of $\widehat{T}^t$ with $A_e \subseteq \widehat{P}^t_{\gamma(t,u)}$ by Lemma 3.3.

We obtain the tree $\overline{T}$ from the disjoint union of the trees $\widehat{T}^t$ for all $t \in V(T)$ by adding the edges $\gamma(t,u)\gamma(u,t)$ for each $tu \in E(T)$. Let $\overline{P}_u$ be $\widehat{P}^t_u$ for the unique $t \in V(T)$ with $u \in V(\widehat{T}^t)$. Then $\overline{\mathcal{T}} := \big(\overline{T}, (\overline{P}_t)_{t\in V(\overline{T})}\big)$ is a tree-decomposition of G since $\mathcal{T}$ is, since the $\widehat{\mathcal{T}}^t$ are tree decompositions of the torsos H_t and since by definition for $e = tu \in E(T)$ we obtain $A_e = \overline{P}_{\gamma(t,u)} \cap \overline{P}_{\gamma(u,t)}$. It is a canonical tree-decomposition, since $\mathcal{T}$ is canonical, since the family $\big(\mathcal{T}^t\big)_{t\in V(T)}$ is canonical, and since the nodes $\gamma(t,u)$ are canonically defined.

Whilst (i) is true by construction, the nodes added in the construction of $\widehat{\mathcal{T}}^t$ are hub nodes by definition, yielding (ii). Furthermore, (iii) and (iv) follow from Remark 3.2 (ii) and from the observation that for all $tu \in E(T)$ the adhesion sets $\overline{P}_{\gamma(t,u)} \cap \overline{P}_{\gamma(u,t)}$ and $P_t \cap P_u$ are equal.

For (v) consider two profiles Q_1, Q_2 of order at most k of a torso H_t efficiently distinguished by $(A,B) \in N(\mathcal{T}^t)$. If (A,B) is also induced by $\widehat{\mathcal{T}}^t$, then we obtain by construction that there is a separation induced by $\overline{\mathcal{T}}$ which induces (A,B) on H_t. If not, then by Remark 3.2 (iv) we obtain that the parts P^t_u and P^t_v for the edge $uv \in E(T^t)$ inducing (A,B) are equal, and since $A \cap B = P^t_u \cap P^t_v = P^t_u$ by definition, we obtain $|P^t_u| < k$. Hence by Corollary 2.6 neither Q_1 nor Q_2 inhabits P^t_u or P^t_v. Let $P^t_{u_i}$ be a part inhabited by Q_i for $i \in \{1,2\}$, which exists by Lemma 2.7. Then the path in T^t from u_1 to u_2 uses the edge uv. If no separation induced by an edge of this path is also induced by an edge of $\widehat{\mathcal{T}}^t$, then all parts corresponding to the nodes on the path are equal to P^t_u, contradicting that Q_i inhabits $P^t_{u_i}$ again with Corollary 2.6. Hence there is a separation (C,D) induced by both an edge of this path and $\widehat{\mathcal{T}}^t$. By (T3) the separator $C \cap D$ is a subset of P^t_u, and therefore (C,D) also distinguishes Q_1 and Q_2 efficiently. □

3.2 Obtaining tree-decompositions from almost nested sets of separations

Theorem 2.3 tells how one can transform a nested set of separations into a tree-decomposition. The main result of this subsection can be seen as a generalization of this theorem.

For a separation (A, B) of G and $X \subseteq V(G)$, the pair $(A \cap X, B \cap X)$ is a separation of $G[X]$, which we call the *restriction* $(A, B)\restriction X$ of (A, B) to X. It may happen that $(A, B)\restriction X$ is improper although (A, B) is proper. The *restriction* $S\restriction X$ to X of a set S of separations of G to X consists of the proper separations $(A, B)\restriction X$ where $(A, B) \in S$.

For a set S of separations of G let $\min_{\text{ord}}(S)$ denote the set of those separations in S with minimal order. Note that if S is nonempty, then so is $\min_{\text{ord}}(S)$, and that $\min_{\text{ord}}$ is canonical in that it commutes with graph isomorphisms.

A finite sequence $(\beta_0, \dots, \beta_n)$ of vertex sets of G is called an *S-focusing sequence* if

(F1) $\beta_0 = V(G)$;

(F2) for all $i < n$, the separation system N_{β_i} generated by $\min_{\text{ord}}(S\restriction\beta_i)$ is nonempty and is nested with the set $S\restriction\beta_i$;

(F3) β_{i+1} is an N_{β_i}-block of $G[\beta_i]$.

An S-focusing sequence $(\beta_0, \dots, \beta_n)$ is *good* if

(F$*$) the separation system N_{β_n} generated by $\min_{\text{ord}}(S\restriction\beta_n)$ is nested with the set $S\restriction\beta_n$.

The set of all S-focusing sequences is partially ordered by extension, where $(V(G))$ is the smallest element. The subset $\mathcal{F}_S$ of all good S-focusing sequences is downwards closed in this partial order. Note that for any S-focusing sequence $(\beta_0, \dots, \beta_n)$ the separation system N_{β_n} is proper.

Lemma 3.5. *Let $(\beta_0, \dots, \beta_n) \in \mathcal{F}_S$ and let $(A, B) \in S$. If $(A, B)\restriction\beta_n$ is proper, then $A \cap B \subseteq \beta_n$.*

Proof. By assumption $(A, B)\restriction\beta_n$ is proper, hence there are $a \in (\beta_n \cap A) \setminus B$ and $b \in (\beta_n \cap B) \setminus A$. Since $\beta_n \subseteq \beta_i$ for all $i \leq n$ the separations $(A, B)\restriction\beta_i$ are proper as well. Suppose for a contradiction there is a vertex $v \in (A \cap B) \setminus \beta_n$. Let $j < n$ be maximal with $v \in \beta_j$. Since β_{j+1} is an N_{β_j}-block of $G[\beta_j]$, there is a separation $(C, D) \in N_{\beta_j}$ with $v \in C \setminus D$ and $\{a, b\} \subseteq \beta_n \subseteq \beta_{j+1} \subseteq D$.

Now a, b and v witness that $(A,B)\restriction\beta_j$ and (C,D) are not nested: Indeed, a witnesses that D is not a subset of $B\cap\beta_j$. Similarly, b witnesses that D is not a subset of $A\cap\beta_j$. But v witnesses that neither $A\cap\beta_j$ nor $B\cap\beta_j$ is a subset of D. Thus we get a contradiction to the assumption that N_{β_j} is nested with the set $S\restriction\beta_j$. □

A set S of separations of G is *almost nested* if all S-focusing sequences are good. In this case the maximal elements of $\mathcal{F}_S$ of the partial order are exactly the S-focusing sequences $(\beta_0,\dots,\beta_n)$ with $N_{\beta_n}=\emptyset$, and hence $S\restriction\beta_n=\emptyset$.

Lemma 3.6. *Let S be an almost nested set of separations of G.*

(i) If $(\beta_0,\dots,\beta_n)\in\mathcal{F}_S$ is maximal, then β_n is an S-block.

(ii) If b is an S-block, there is a maximal $(\beta_0,\dots,\beta_n)\in\mathcal{F}_S$ with $\beta_n=b$.

Proof. Let $(\beta_0,\dots,\beta_n)\in\mathcal{F}_S$ be maximal. Then $S\restriction\beta_n$ is empty, i.e. no $(A,B)\in S$ induces a proper separation of $G[\beta_n]$. Hence β_n is S-inseparable. For every $v\in V(G)\setminus\beta_n$ there is an $i<n$ and a separation in N_{β_i} separating v from β_n. Hence β_n is an S-block.

Conversely given an S-block b, let $(\beta_0,\dots,\beta_n)\in\mathcal{F}_S$ be maximal with the property $b\subseteq\beta_n$, which exists since $(V(G))\in\mathcal{F}_S$ and since $\mathcal{F}_S$ is finite. Since b is N_{β_n}-inseparable, there is some N_{β_n}-block β_{n+1} containing b. The choice of $(\beta_0,\dots,\beta_n)$ implies that $(\beta_0,\dots,\beta_{n+1})\notin\mathcal{F}_S$ and hence $N_{\beta_n}=\emptyset$, i.e. $(\beta_0,\dots,\beta_n)$ is a maximal element of $\mathcal{F}_S$. Thus β_n is an S-block with $b\subseteq\beta_n$ and hence $b=\beta_n$. □

The *rank* of an S-focusing sequence $(\beta_0,\dots,\beta_n)$ is either the size of any separator $A\cap B$ if there is a separation $(A,B)\in N_{\beta_n}$, or ∞ otherwise.

For an almost nested set S of separations of G two S-focusing sequences $(\beta_0,\dots,\beta_n)$ and $(\alpha_0,\dots\alpha_m)$ are *similar* if there is an automorphism ψ of G inducing an isomorphism between $G[\beta_n]$ and $G[\alpha_m]$. Similar S-focusing sequences clearly have the same rank. If S is canonical, then ψ induces an isomorphism between $\mathcal{T}(N_{\beta_n})$ and $\mathcal{T}(N_{\alpha_m})$ as obtained from Theorem 2.3.

Theorem 3.7. *Let S be a canonical almost nested set of separations of G. Then there is a canonical tree-decomposition $\mathcal{T}$ of G such that*

(i) every S-block of G is a part of $\mathcal{T}$;

(ii) every part of $\mathcal{T}$ is either an S-block of G or a hub;

(iii) and for every separation (A, B) induced by $\mathcal{T}$ there is a separation $(A', B') \in S$ such that $A \cap B = A' \cap B'$.

Proof. We recursively construct for every S-focusing sequence $(\beta_0, \dots, \beta_n)$ a tree-decomposition $\mathcal{T}^{\beta_n}$ of $G[\beta_n]$ so that the tree-decomposition $\mathcal{T}^{V(G)}$ for the smallest S-focusing sequence is the desired tree-decomposition of G. We show inductively that for any S-focusing sequence $(\beta_0, \dots, \beta_n)$ the tree-decomposition $\mathcal{T}^{\beta_n}$ has the following properties:

(a) every S-block included in β_n is a part of $\mathcal{T}^{\beta_n}$;

(b) every part of $\mathcal{T}^{\beta_n}$ is either an S-block or a hub;

(c) every separation (A, B) induced by $\mathcal{T}^{\beta_n}$ is proper;

(d) and for every separation (A, B) induced by $\mathcal{T}^{\beta_n}$ there is a separation $(A', B') \in S$ and an S-focusing sequence $(\beta_0, \dots, \beta) \geq (\beta_0, \dots, \beta_n)$[7] such that $(A', B') \restriction \beta = (A, B)$.

Furthermore we show by induction, that

(e) if $(\alpha_0, \dots, \alpha_n)$ and $(\beta_0, \dots, \beta_n)$ are similar, then $\mathcal{T}^{\alpha_m}$ and $\mathcal{T}^{\beta_n}$ are isomophic.

For the maximal S-focusing sequences we just take the trivial tree-decompositions with only a single part. These tree-decompositions satisfy (a) and (b) by Lemma 3.6, and (c) and (d) since their trees do not have any edges. If for two S-blocks b_1 and b_2 there is an isomorphism between $G[b_1]$ and $G[b_2]$ induced by an automorphism of G, then clearly the tree-decompositions are isomorphic. Hence (e) holds for all S-focusing sequences of rank ∞.

Suppose for our induction hypothesis that for every S-focusing sequence $(\alpha_0, \dots, \alpha_m)$ with rank greater than r there is a canonical tree-decomposition $\mathcal{T}^{\alpha_m}$ of $G[\alpha_m]$ with the desired properties. Furthermore suppose that that (e) holds for all $(\alpha_0, \dots, \alpha_m) \in \mathcal{F}_S$ with rank greater than r.

To construct $\mathcal{T}^{\beta_n}$ for an S-focusing sequence $(\beta_0, \dots, \beta_n)$ of rank r, first we take the tree-decomposition $\mathcal{T}(N_{\beta_n})$ of $G[\beta_n]$ as in Theorem 2.3. Next we define a canonical tree-decomposition of each torso H_t of $\mathcal{T}(N_{\beta_n})$. If t is a hub node, we take the trivial tree-decomposition. Otherwise P_t is an N_{β_n}-block β by Theorem 2.3 (*ii*) and we take $\mathcal{T}^{\beta}$. This is indeed a tree-decomposition of H_t: for a separation (A, B) induced by $\mathcal{T}^{\beta}$ consider (A', B')

[7] Recall that the order on $\mathcal{F}_S$ is given by extension.

as given in (d). By Lemma 3.5 we obtain that $A' \cap B' \subseteq \beta$ and hence $A' \cap B' = A \cap B$. Therefore if (A, B) would separate an adhesion set $C \cap D$ for $(C, D) \in N_{\beta_n}$, then so would $(A', B') \restriction \beta_n$, contradicting (F$*$).

If two torsos H_t and H_u of $\mathcal{T}(N_{\beta_n})$ are similar, then either $V(H_t)$ and $V(H_u)$ are $N(\beta_n)$-blocks whose corrensponding S-focusing sequences have rank greater than r, or they are hubs. If they are N_{β_n}-blocks, the chosen tree-decompositions are isomorphic by the induction hypothesis. If they are hubs, the chosen trivial tree-decompositions are isomorphic as witnessed by the same automorphism of G witnessing the similarity of H_t and H_u. Hence this family of tree-decompositions of the torsos of $\mathcal{T}(N_{\beta_n})$ is canonical.

We apply Lemma 3.4 to $\mathcal{T}(N_{\beta_n})$ and the family of tree-decompositions of the torsos to get a canonical tree-decomposition $\mathcal{T}^{\beta_n}$ of $G[\beta_n]$ for $(\beta_0, \dots, \beta_n)$, which satisfies (a), (b) and (c) by Lemma 3.4 (i), (ii) and (iii) and by the induction hypothesis. Also by Lemma 3.4 (iii) for a separation (A, B) induced by $\mathcal{T}^{\beta_n}$ either $(A, B) \in N_{\beta_n} \subseteq S \restriction \beta_n$ or (A, B) induces a separation in $\mathcal{T}^{\beta}$ for an N_{β_n}-block β on the corresponding torso. In the first case $(\beta_0, \dots, \beta_n)$ is the desired S-focusing sequence for (d) and in the second case the induction hypothesis yields $(A', B') \in S$ and the desired S-focusing sequence extending $(\beta_0, \dots, \beta_n, \beta)$. Hence (d) holds for $\mathcal{T}^{\beta_n}$.

Let $(\alpha_0, \dots, \alpha_m)$ be similar to $(\beta_0, \dots, \beta_n)$. Since S is canonical, the automorphism of G that witnesses the similarity also witnesses that $\mathcal{T}(N_{\alpha_m})$ and $\mathcal{T}(N_{\beta_n})$ are isomorphic. Hence any torso of $\mathcal{T}(N_{\alpha_m})$ is similar to the corrensponding torso of $\mathcal{T}(N_{\beta_n})$ and by induction hypothesis the tree-decompositions of the torsos are isomorphic. Therefore following the construction of Lemma 3.4 yields (e).

Inductively the tree-decomposition $\mathcal{T}^{V(G)}$ of G is canonical and satisfies (i) and (ii) by (a) and (b). Finally, (iii) follows from (c), (d) and Lemma 3.5. $\square$

3.3 Extending a nested set of separations

Let $\mathcal{T} := \big(T, (P_t)_{t \in V(T)}\big)$ be a tree-decomposition and let $N := N(\mathcal{T})$ be the induced nested set of separations. Recall that a separation (A, B) of G that is nested with N *induces* a separation $(A \cap P_t, B \cap P_t)$ of the torso H_t. A k-profile $\widetilde{Q}$ of H_t is *induced* by a k-profile Q of G if for every $(A', B') \in \widetilde{Q}$ there is an $(A, B) \in Q$ which induces (A', B') on H_t.

Lemma 3.8. *Let $t \in V(T)$ and for $i \in \{1, 2\}$ let Q_i be a k_i-profile of G inhabiting P_t.*

(i) Q_i induces a unique k_i-profile $\widetilde{Q}_i$ of H_t;

(ii) *if a separation* (A, B) *of* G *nested with* N *distinguishes* Q_1 *and* Q_2 *(efficiently), then the by* (A, B) *induced separation on* H_t *distinguishes the induced profiles* $\widetilde{Q}_1$ *and* $\widetilde{Q}_2$ *of* H_t *(efficiently);*

(iii) *if a separation* (A, B) *of* H_t *distinguishes the induced profiles* $\widetilde{Q}_1$ *and* $\widetilde{Q}_2$ *of* H_t *(efficiently), then any separation of* G *nested with* N *that induces* (A, B) *on* H_t *distinguishes* Q_1 *and* Q_2 *(efficiently);*

(iv) *if* Q_i *is* $(k-1)$*-robust, then so is* $\widetilde{Q}_i$*.*

Proof. Let (A, B) be a proper separation of order less than k_i of H_t. Applying Lemma 2.5 to $X := A \cap B$ yields a component C of $G - X$ such that $(V(G) \setminus C, C \cup X) \in Q_i$. Since Q_i inhabits P_t, this separation witnesses that $C \cap P_t$ is not empty. By construction $C \cap P_t$ is a subset of either $A \setminus B$ or $B \setminus A$, say $B \setminus A$. For every $v \in P_t \setminus X$ let C_v be the component of $G - X$ with $v \in C_v$. Let Y be the union of all components C_v with $v \in B \setminus A$. Then the separation $(V(G) \setminus Y, X \cup Y)$ is in Q_i by consistency and induces (A, B).

Hence we obtain an orientation $\widetilde{Q}_i$ of separations of order less than k_i of H_t by additionally taking every improper separation $(A, V(G))$ of order less than k_i to be in $\widetilde{Q}_i$. Consistency and (P) for this orientation are inherited from Q_i, hence $\widetilde{Q}_i$ is a k_i-profile of H_t induced by Q_i.

Let (A, B) be a separation of G nested with N distinguishing Q_1 and Q_2 (efficiently). Since both Q_1 and Q_2 inhabit P_t we obtain that $(A \setminus B) \cap P_t$ and $(B \setminus A) \cap P_t$ are both not empty. Therefore the separation that (A, B) induces on H_t is proper and hence distinguishes $\widetilde{Q}_1$ and $\widetilde{Q}_2$ (efficiently).

If a separation (A, B) of G nested with N does not distinguish Q_1 and Q_2 (efficiently), then the induced separation does not distinguish $\widetilde{Q}_1$ and $\widetilde{Q}_2$ (efficiently) as well.

Hence *(ii)* and *(iii)* hold, implying the uniqueness of *(i)*. *(iv)* holds since the corner separations of two induced separations equal the induced separations of the corresponding corner separations of the original separations. □

For a nested separation system N let $S^N_{<k}$ be the set of separations of order less than k of G nested with N.

Theorem 3.9. *Let* N *be a canonical nested system of separations of order less than* k *of* G *such that* $S^N_{<k}$ *distinguishes any two* k*-distinguishable* $(k-1)$*-robust profiles of order at most* k *of* G *efficiently. Then there is a canonical nested system* $\widehat{N} \supseteq N$ *of separations of order less than* k *of* G *distinguishing any two* k*-distinguishable* $(k-1)$*-robust profiles of order at most* k *of* G *efficiently.*

Proof. Let $\mathcal{T}(N) = \big(T, (P_t)_{t\in V(T)}\big)$ be the canonical tree-decomposition obtained with Theorem 2.3 from N. We define a canonical family $\big(\mathcal{T}^t\big)_{t\in V(T)}$ of tree-decompositions of the torsos recursively. For any $t \in V(T)$ for which $\mathcal{T}^t$ has not been defined let $\mathcal{T}^t$ be a canonical tree-decomposition of adhesion less than k of the torso H_t that distinguishes all the k-distinguishable $(k-1)$-robust profiles of order $\leq k$ of H_t efficiently which exists by Theorem 2.8. For every torso H_u similar to H_t witnessed by an automorphism ψ of G we then define $\mathcal{T}^u$ to be the tree-decomposition of H_u isomorphic to $\mathcal{T}^t$ as witnessed by ψ. We repeat this procedure until family $\big(\mathcal{T}^t\big)_{t\in V(T)}$ has been defined.

Let N^t denote the canonical nested separation system induced by $\mathcal{T}^t$. By applying Lemma 3.4 to $\mathcal{T}(N)$ and $\big(\mathcal{T}^t\big)_{t\in V(T)}$ we obtain a canonical tree-decomposition $\widehat{\mathcal{T}}$ of G and we claim that the canonical nested separation system $\widehat{N}$ induced by this tree-decomposition has the desired properties by construction and Lemma 3.8.

We obtain $\widehat{N} \supseteq N$ by Lemma 3.4 (iv). Let Q_1 and Q_2 be any two k-distinguishable $(k-1)$-robust profiles of order at most k of G. Suppose Q_1 and Q_2 are not already efficiently distinguished by N. Let $(A,B) \in S^N_{<k}$ distinguish Q_1 and Q_2 efficiently and let P_t be a part of $\mathcal{T}(N)$ such that $A \cap B \subseteq P_t$.

If Q_1 and Q_2 both inhabit P_t, then there is a separation $(C,D) \in N^t$ distinguishing the induced profiles $\widetilde{Q}_1$ and $\widetilde{Q}_2$ efficiently. By Lemma 3.4 (v) and Lemma 3.8 (iii) there is a separation in $\widehat{N}$ distinguishing Q_1 and Q_2 efficiently.

If for both $i \in \{1,2\}$ we obtain that Q_i inhabits $P_{t_i} \neq P_t$, then consider the neighbour u_i of t on the path between the t and t_i. Let k_i denote the size of the adhesion set $P_t \cap P_{u_i}$ and let b_i denote the k_i-block of H_t containing $P_t \cap P_{u_i}$. Since (A,B) distinguishes Q_1 and Q_2, we obtain that b_1 and b_2 lie on different sides of (A,B). By the assumption that N does not distinguish Q_1 and Q_2 efficiently, we obtain that the order of (A,B) is less than k_i, and hence $b_1 \neq b_2$. The induced profiles $P_{k_1}(b_1)$ and $P_{k_2}(b_2)$ are both $(k-1)$-robust by Lemma 2.4 and k-distinguishable, since by Lemma 3.8 the separation that (A,B) induces on H_t distinguishes them. Again with Lemma 3.4 (v) we obtain a separation $(C,D) \in \widehat{N}$ inducing a separation on H_t that distinguishes $P_{k_1}(b_1)$ and $P_{k_2}(b_2)$ efficiently. Since by construction P_{t_1} and P_{t_2} lie on different sides of (C,D), we obtain that (C,D) distinguishes Q_1 and Q_2 efficiently.

If only one Q_i inhabits P_t we obtain the theorem with a combination of both of the above arguments. □

4 Proof of the main result

Recall that a k-block b is *polishable* if there is a tree-decomposition of adhesion less than k of G in which b is a part. Given a k-block b, let $S(b)$ be the set of all tight separations (A, B) with $b \subseteq B$ such that $A \setminus B$ is a component of $G - b$. Note that $S(b)$ is a nested set of separations, while for different k-blocks b, b' the union $S(b) \cup S(b')$ need not to be nested.

Lemma 4.1. *Let b be a k-block of G. If b is polishable, then all separations in $S(b)$ have order less than k.*

Proof. Let $\mathcal{T} = \big(T, (P_t)_{t \in V(T)}\big)$ be a tree-decomposition of adhesion less than k of G with $P_t = b$ for some $t \in V(T)$. Each separation in $S(b)$ has the form $(K \cup N(K), V(G) \setminus K)$ for some component K of $G - b$. There is a neighbour u of t such that the separation (A, B) induced by tu satisfies $b \subseteq B$ and $K \subseteq A \setminus B$. As $N(K) \subseteq A \cap B$, we have $|N(K)| < k$, completing the proof. □

Hence for $S(b)$ to have only separations of order less than k is a necessary condition for a k-block b to be polishable. Theorem 4.8 will imply that this condition is also sufficient.

Remark 4.2. *Let b be a k-block of G. For all $(A, B) \in S(b)$ the separator $A \cap B$ is a subset of b.*

Proof. Since $A \setminus B$ is a component of $G - b$, the neighourhood of $A \setminus B$ in G is a subset of b. And since (A, B) is tight, the neighbourhood of $A \setminus B$ in G is $A \cap B$. □

For the remainder of this section we will focus on the following canonical set of separarions of order less than k.

$$S := \bigcup \big\{ S(b) \cap S_{<k} \;\big|\; b \text{ is a } k\text{-block of } G \big\}.$$

Lemma 4.3. *Every k-block b of G with $S(b) \subseteq S_{<k}$ is an S-block.*

Proof. Clearly, any k-block b is S-inseparable. For $v \in V(G) \setminus b$ let C be the component of $G - b$ with $v \in C$. Every vertex w in the neighbourhood of C also has a neighbour in $V(G) \setminus (C \cup N(C))$ since otherwise it could be separated from b by $N(C) \setminus \{w\}$. Hence the separation $(C \cup N(C), V(G) \setminus C)$ is tight and therefore in $S(b) \subseteq S$ by construction. Hence b is an S-block. □

Lemma 4.4. *Let $(A, B) \in S$ and (C, D) an arbitrary separation of G. If the link ℓ_A is empty, then (A, B) and (C, D) are nested.*

Proof. Since $A \setminus B$ is connected, either $\mathrm{int}(A, C)$ or $\mathrm{int}(A, D)$ is empty, say $\mathrm{int}(A, C)$. Since (A, B) is tight, the link ℓ_C is empty. Hence by Remark 2.1 (A, B) and (C, D) are nested. □

Lemma 4.5. *Let $(A, B), (C, D) \in S$ be crossing. Then the links ℓ_B and ℓ_D are empty and the corner $B \cap D$ is the union of left sides E of separations $(E, F) \in S$, all of whose orders are strictly smaller than the orders of both (A, B) and (C, D).*

Proof. Let b_1 and b_2 be k-blocks with $(A, B) \in S(b_1)$ and $(C, D) \in S(b_2)$. By Lemma 4.4, ℓ_A and ℓ_C are not empty. Since by Remark 4.2, the separator $A \cap B$ is included in b_1 and since (C, D) cannot separate b_1, the link ℓ_D is empty. Since b_1 includes the nonempty link ℓ_C, we obtain $b_1 \subseteq C$ and hence $b_1 \subseteq B \cap C$. Similarly we get that ℓ_B is empty and $b_2 \subseteq A \cap D$.

Let K be an arbitrary component of $G[\mathrm{int}(B, D)]$. Let $E := K \cup N(K)$ and $F := V(G) \setminus K$. Since the center c is a subset of $b_1 \cap b_2$ and since $K \cap (b_1 \cup b_2)$ is empty, K is a component of $G - b_1$. Since every vertex in $E \cap F \subseteq c$ also has a neighbour in $A \setminus B \subseteq F \setminus E$, we obtain $(E, F) \in S(b_1)$. And since $E \cap F \subseteq c$ and since ℓ_A and ℓ_C are not empty, we obtain $|E \cap F| < \min\{|A \cap B|, |C \cap D|\}$. □

Lemma 4.6. *S is almost nested.*

Proof. We have to show that every S-focusing sequence $(\beta_0, \ldots, \beta_n)$ is good, i.e. N_{β_n} is nested with $S \restriction \beta_n$. Let $(\beta_0, \ldots, \beta_n)$ be an S-focusing sequence. Let $(A, B) \restriction \beta_n \in N_{\beta_n}$ and $(C, D) \restriction \beta_n \in S \restriction \beta_n$. If (A, B) and (C, D) are nested, then so are $(A, B) \restriction \beta_n$ and $(C, D) \restriction \beta_n$. Suppose (A, B) and (C, D) are crossing.

By Lemma 4.5 ℓ_B and ℓ_D is empty. If $\mathrm{int}(B, D) \cap \beta_n$ is empty, then by Remark 2.1 $(A, B) \restriction \beta_n$ and $(C, D) \restriction \beta_n$ are nested. Hence by Lemma 4.5 it suffices to show that $(E \setminus F) \cap \beta_n$ is empty for every $(E, F) \in S$ with $E \subseteq B \cap D$ whose order is strictly smaller than the order of (A, B).

Since $(A, B) \restriction \beta_n$ has minimal order among all separations in $S \restriction \beta_n$, we obtain that $(E, F) \restriction \beta_n$ is improper. As $A \subseteq F$, the set $(F \setminus E) \cap \beta_n$ is not empty, thus $(E \setminus F) \cap \beta_n$ is empty, as desired. □

By combining Theorem 3.7 with Lemma 4.3 and Lemma 4.6 we obtain a canonical tree-decomposition $\widehat{\mathcal{T}}$ of G in which every k-block b with $S(b) \subseteq S_{<k}$ is a part of $\widehat{\mathcal{T}}$. Let N be the canonical nested separation system induced by $\widehat{\mathcal{T}}$ and like before $S^N_{<k}$ the set of separations of order less than k of G nested with N.

Lemma 4.7. *$S_{<k}^N$ distinguishes any two k-distinguishable $(k-1)$-robust profiles of order at most k of G efficiently.*

Proof. Let Q_1 and Q_2 be k-distinguishable $(k-1)$-robust profiles of order at most k of G. Let $(A, B) \in Q_1$ distinguish Q_1 and Q_2 efficiently such that the (finite) cardinality of the set of separations $(C, D) \in N$ that cross (A, B) is minimal.

Suppose for a contradiction that there is a separation $(C, D) \in N$ that crosses (A, B). Since by Theorem 3.7 (iii) the separator $C \cap D$ coincides with the separator of a separation in S, Remark 4.2 implies that $C \cap D$ is $S_{<k}$-inseparable and hence either ℓ_A or ℓ_B is empty. Without loss of generality let ℓ_B be empty. The order of the corner separations $(A \cup D, B \cap C)$ and $(A \cup C, B \cap D)$ is less or equal than $|A \cap B|$, hence they are oriented by Q_1 and Q_2. Applying Lemma 2.5 to $X := A \cap B$ and Q_1 yields a component K of $G - X$ with $(V(G) \setminus K, K \cup X) \in Q_1$. In particular we get $K \subseteq B \setminus A$ by consistency. Since ℓ_B is empty and K is connected, we obtain $K \subseteq C \setminus D$ or $K \subseteq D \setminus C$. Therefore either $(A \cup D, B \cap C)$ or $(A \cup C, B \cap D)$ is in Q_1 by consistency to $(V(G) \setminus K, K \cup X)$, and not in Q_2 by consistency to (B, A).

Hence there is a corner separation of (A, B) and (C, D) distinguishing Q_1 and Q_2 efficiently. By Lemma 2.2 it is nested with every separation in N that is also nested with (A, B), as well as with (A, B) itself. Hence it crosses less separations of N than (A, B), contradicting the choice of (A, B). Thus (A, B) is nested with N. □

Theorem 4.8. *Every graph G has a canonical tree-decomposition $\mathcal{T}$ of adhesion less than k distinguishing any two k-distinguishable $(k-1)$-robust profiles of order at most k of G efficiently such that every k-block b of G with $S(b) \subseteq S_{<k}$ is equal to the unique part of $\mathcal{T}$ in which it is contained.*

Proof. By Theorem 3.9 and Lemma 4.7 the canonical nested separation system N as before can be refined to a set $\widehat{N}$ that distinguishes all k-distinguishable $(k-1)$-robust profiles of order at most k of G efficiently and hence $\mathcal{T} := \mathcal{T}(\widehat{N})$ as in Theorem 2.3 is the tree-decomposition with the desired properties. □

With Lemma 4.1 this yields the characterization in Corollary 1.2. Hence we can restate this theorem in terms of polishable k-blocks.

Corollary 4.9. *Every graph G has a canonical tree-decomposition $\mathcal{T}$ of adhesion less than k distinguishing any two k-distinguishable $(k-1)$-robust profiles of order at most k of G efficiently such that every polishable k-block of G is equal to the unique part of $\mathcal{T}$ in which it is contained.* □

5 Comparison with other tree-decompositions

Every separation that is induced by any of the tree-decompositions constructed in [3] is *essential*, in that it distinguishes two k-profiles efficiently.

Example 5.1. This example shows a graph where the tree-decompositions constructed in [3] do not bring out the polishable 4-blocks. Consider the graph obtained from two disjoint cliques on four vertices and another vertex v by connecting v to two vertices of each clique, as depicted in Figure 4. The only essential separation is depicted in black and hence is the only separation induced by an edge of any tree-decomposition constructed in [3]. But both 4-blocks are polishable, hence Theorem 4.8 yields the tree-decomposition depicted in gray.

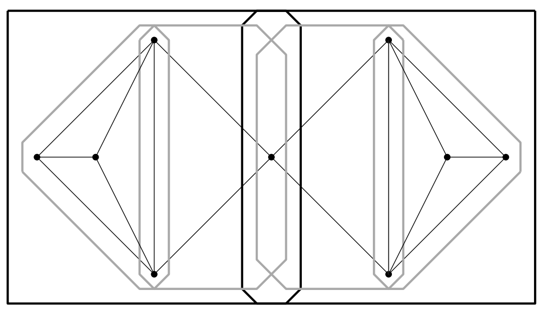

Figure 4: two canonical tree-decompositions

A tree-decomposition $\mathcal{T} := \big(T, (P_t)_{t\in V(T)}\big)$ is *lean* if for any two nodes $t_1, t_2 \in V(T)$ and vertex sets $X_1 \subseteq P_{t_1}$ and $X_2 \subseteq P_{t_2}$ with $|X_1| = |X_2| =: k$, either G contains k disjoint $X_1 - X_2$ paths or an adhesion set along the path $t_1 T t_2$ has size less than k.

Proposition 5.2. *Every polishable k-block b appears as a part in any lean tree-decomposition of adhesion less than k of a graph G.*

Proof. Let $\mathcal{T}$ be a lean tree-decomposition of adhesion less than k and let P_t be the part of $\mathcal{T}$ with $b \subseteq P_t$. Suppose for a contradiction that there is a vertex $v \in P_t \setminus b$. Using Lemma 4.1 we obtain a tight separation $(A, B) \in S(b)$ of order less than k with $v \in A \setminus B$ and $b \subseteq B$. With Remark 4.2 we obtain that $A \cap B \subsetneq b \subsetneq P_t$. Let $w \in b \setminus (A \cap B)$, let $X_1 := (A \cap B) \cup \{w\}$ and let $X_2 := (A \cap B) \cup \{v\}$. Since $A \cap B$ separates v from w, there is no $v - w$ path avoiding $A \cap B$ and hence no $|A \cap B| + 1$ disjoint $X_1 - X_2$ paths, contradicting the leanness of $\mathcal{T}$. □

6 Appendix

In this appendix we will relate the notion of almost nested sets of separations with the notion of trees of tree-decompositions given in [2] to obtain Theorem 2.8. This also yields an alternative characterization for almost nested sets of separations in the process.

A *tree of tree-decompositions* $\mathcal{U} = \big(U, (\mathcal{T}^u)_{u \in V(U)}\big)$ of a graph G consists of a rooted tree U and for every node $u \in V(U)$ a graph H^u and a tree-decomposition $\mathcal{T}^u = \big(T^u, (P^u_t)_{t \in V(T^u)}\big)$ of H^u, such that

(a) $H^r = G$ for the root r of U;

(b) the graphs assigned to the children[8] of $u \in V(U)$ are the torsos of the parts of $\mathcal{T}^u$;

(c) if $u \in V(U)$ is at level[9] k, then every adhesion set of $\mathcal{T}^u$ has size k.

A tree of tree-decompositions is *canonical* if each tree-decomposition $\mathcal{T}^u$ for $u \in V(U)$ is canonical and if for every pair of nodes where the assigned graphs H^u and H^v are similar[10] the corresponding tree-decompositions are isomorphic as witnessed by the same automorphism of G that witnesses the similarity of H^u and H^v.

A part P^u_t for $u \in V(U)$ and $t \in V(T^u)$ is called *final* if it is not a hub and for every descendant[11] u' of u the tree-decomposition $\mathcal{T}^{u'}$ is trivial.

Proposition 6.1. *For every canonical almost nested set S of separations there is a canonical tree of tree-decompositions $\mathcal{U}$ such that*

(i) the S-blocks are precisely the final parts of $\mathcal{U}$;

(ii) for every separation $(A, B) \in S$ there is a node $u \in V(U)$ and a separation $(A', B') \in N(\mathcal{T}^u)$ such that (A, B) induces (A', B') on H^u.

Proof. For every S-focusing sequence $(\beta_0, \ldots, \beta_n) \in \mathcal{F}_S$ we obtain a canonical tree-decomposition $\mathcal{T}(N_{\beta_n})$ of $G[\beta_n]$ by Theorem 2.3. As $N_{\beta_{n-1}}$ is nested with $S{\restriction}\beta_{n-1}$, in particular with the separations that induce N_{β_n} on $G[\beta_n]$, no separation in N_{β_n} separates an adhesion set and thus $\mathcal{T}(N_{\beta_n})$ is a tree-decomposition of the torso H_t. By construction all adhesion sets of the same

[8] A *child* of a node u is a neighbour with greater distance to the root.

[9] The k-th *level* is the set of all nodes with distance k to the root.

[10] As before we call H^u and H^v *similar* if there is an automorphism of G inducing an isomorphism between H^u and H^v.

[11] The *descendance* of a node u are its children combined with their descendance.

tree-decomposition have the same size. Hence we obtain $\mathcal{U}$ by adding trivial tree-decompositions of the torsos if the difference in this size between two neighbouring S-focusing sequences is greater than one. With Lemma 3.6 we obtain (i) and we obtain (ii) by construction. By construction $\mathcal{U}$ is canonical since S is. □

Lemma 6.2. *Let $\mathcal{T}$ be a tree-decomposition of G, let H_t be a torso for a node t and let (A, B) be a proper separation of H_t. Then there is a separation of G nested with $N(\mathcal{T})$ that induces (A, B) on H_t.*

Proof. As in the proof of Lemma 3.8 we obtain that there is a separation of G that induces (A, B) and has the same separator. Let (A', B') be a separation of G inducing (A, B) on H_t with $A \cap B = A' \cap B'$ such that the number of separations in $N(\mathcal{T})$ that cross (A', B') is minimal. Suppose for a contradiction that there is a separation $(C, D) \in N(\mathcal{T})$ with $V(H_t) \subseteq C$ that crosses (A', B'). The link ℓ_D is a subset of $D \setminus C$, and since $A' \cap B' \subseteq V(H_t)$, we obtain that ℓ_D is empty. Since (A', B') does not separate $C \cap D$, we obtain that either $\ell_{A'}$ or $\ell_{B'}$ is empty, say $\ell_{B'} = \emptyset$. Then by Lemma 2.2 the corner separation $(A' \cup C, B' \cap D)$ is nested with (C, D). It still induces (A, B) on H_t by construction, and it has the separator $A \cap B$, contradicting the choice of (A', B'). □

Proposition 6.3. *For every canonical tree of tree-decompositions $\mathcal{U}$ there is a canonical almost nested set S of separations such that*

(i) the S-blocks are precisely the final parts of $\mathcal{U}$;

(ii) for every separation $(A', B') \in N(\mathcal{T}^u)$ for any $u \in V(U)$ there is a separation $(A, B) \in S$ inducing (A', B') on H^u.

Proof. Let $\big(U, (\mathcal{T}^u)_{u \in V(U)}\big) := \mathcal{U}$. Let S be the set of all separations (A, B) of G such that

(a) for all $u \in V(U)$ the separation $(A, B) \restriction V(H^u)$ is either improper or nested with $N(\mathcal{T}^u)$;

(b) there is a node $u \in V(U)$ and a separation $(A', B') \in N(\mathcal{T}^u)$ such that $(A, B) \restriction V(H^u) = (A', B')$.

By (b), S is canonical since $\mathcal{U}$ is. If for every node $u \in V(U)$ and every $(A', B') \in N(\mathcal{T}^u)$ there is an $(A, B) \in S$ such that $(A, B) \restriction V(H^u) = (A', B')$, then S is almost nested by (a) and the S-blocks are the final parts by construction.

Let $u \in V(U)$ and $(A', B') \in N(\mathcal{T}^u)$. Applying Lemma 6.2 successively yields a separation (A, B) of G such that for all s along the path P between the root of U and u we obtain $(A, B)\restriction V(H^s)$ is nested with $N(\mathcal{T}^s)$ and $(A, B)\restriction V(H^u) = (A', B')$. For every node s not on P we obtain that $(A, B)\restriction V(H^u)$ is improper by construction. Hence we obntain $(A, B) \in S$, as required. □

Let $\mathcal{U} := \big(U, (\mathcal{T}^u)_{u \in V(U)}\big)$ be a tree of tree-decompositions of G and let Q be a k-profile of G. By applying Lemma 3.8 successively we obtain for every $1 \leq \ell \leq k$ a unique node $u_\ell \in V(U)$ on level ℓ such that the from Q induced k-profile on $H^{u_{\ell-1}}$ inhabits the part $V(H^{u_\ell})$ of $\mathcal{T}^{u_{\ell-1}}$ tree-decomposition. With setting u_0 to be the root of U, we obtain that u_ℓ is a child of $u_{\ell-1}$. We call the path $u_0 \dots u_k$ the *induced path of Q in U*.

A k_1-profile Q_1 and a k_2-profile Q_2 are *distinguished* by $\mathcal{U}$ if for some $\ell \leq \min\{k_1, k_2\}$ the the induced paths differ somewhere on the first $\ell + 1$ nodes. They are distinguished *efficiently* if the last common node of the induced paths is at level $|A \cap B|$ for any separation (A, B) that efficiently distinguishes Q_1 and Q_2.

Hence Proposition 6.1 (*ii*), Proposition 6.3 (*ii*) and Lemma 3.8 yield

Corollary 6.4. *A tree of tree-decompositions distinguishes two k-profiles efficiently if and only if the corresponding almost nested set of separations does.* □

With this characterization, Theorem 2.8 follows from [2, Theorem 9.2].

References

[1] B. Behrens. Weiteparameter in der Minorentheorie endlicher Graphen. Bachelor's thesis, University of Hamburg, 2014.

[2] J. Carmesin. On the end structure of infinite graphs. arXiv:1409.6640, 2014.

[3] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs I. Existence and algorithms. *J. Combin. Theory (Series B)*, (to appear).

[4] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark. Canonical tree-decompositions of finite graphs II. The parts. *J. Combin. Theory (Series B)*, (to appear).

[5] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. Connectivity and tree structure in finite graphs. *Combinatorica*, 34:1146, 2014.

[6] J. Carmesin and J.P. Gollin. Polishing tree-decompositions to bring out the k-blocks. Preprint 2014.

[7] R. Diestel. Open problems about canonical tree-decompositions of graphs and width parameter duality. Presentation at *Hamburg workshop on graphs and matroids*, Spiekeroog 2014.

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

[9] F. Hundertmark and S. Lemanczyk. Profiles. An algebraic approach to combinatorial connectivity. In preparation.

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

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

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