

k -Blocks: a connectivity invariant for graphs

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

June 25, 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 fewer than k other vertices. The *block number* $\beta(G)$ of G is the largest integer k such that G has a k -block.

We investigate how β interacts with density invariants of graphs, such as their minimum or average degree. We further present algorithms that decide whether a graph has a k -block, or which find all its k -blocks.

The connectivity invariant $\beta(G)$ has a dual width invariant, the *block-width* $\text{bw}(G)$ of G . Our algorithms imply the duality theorem $\beta = \text{bw}$: a graph has a *block-decomposition* of width and adhesion $< k$ if and only if it contains no k -block.

1 Introduction

Given $k \in \mathbb{N}$, a set I of at least k vertices of a graph G is $(< k)$ -inseparable if no set S of fewer than k vertices of G separates any two vertices of $I \setminus S$ in G . A maximal $(< k)$ -inseparable set is a k -block. The degree of connectedness of such a set of vertices is thus measured in the ambient graph G , not only in the subgraph they induce. While the vertex set of a k -connected subgraph of G is clearly $(< k)$ -inseparable in G , there can also be k -blocks that induce few or no edges.

The k -blocks of a graph were first studied by Mader [12]. They have recently received some attention because, unlike its k -connected subgraphs, they offer a meaningful notion of the ‘ k -connected pieces’ into which the graph may be decomposed [2]. This notion is related to, but not the same as, the notion of a tangle in the sense of Robertson and Seymour [14]; see Section 6 and [7] for more on this relationship.

Although Mader [10] had already proved that graphs of average degree at least $4(k-1)$ have k -connected subgraphs, and hence contain a k -block, he did not in [12] consider the analogous extremal problem for the weaker notion of a k -block directly.

Our aim in this paper is to study this problem: we ask what average or minimum degree conditions force a given finite graph to contain a k -block.

This question can, and perhaps should, be seen in a wider extremal context. Let $\beta(G)$ denote the *block number* of G , the greatest integer k such that G has

a k -block (equivalently: has a $(< k)$ -inseparable set of vertices). This β seems to be an interesting graph invariant¹, and one may ask how it interacts with other graph invariants, not just the average or minimum degree. Indeed, the examples we describe in Section 3 will show that containing a k -block for large k is compatible with having bounded minimum and average degree, even in all subgraphs. So k -blocks can occur in very sparse graphs, and one will need bounds on other graph invariants than δ and d to force k -blocks in such graphs.

There is an invariant dual to β : the least integer k such that a graph G has a *block-decomposition* of adhesion and width both at most k . Calling this k the *block-width* $\text{bw}(G)$ of G , we can express the duality neatly as $\beta = \text{bw}$.

All the graphs we consider are finite. Our paper is organized as follows. In Section 2 we introduce whatever terminology is not covered in [3], and give some background on tree-decompositions. In Section 3 we present examples of k -blocks, aiming to exhibit the diversity of the concept. In Section 4 we prove that graphs of minimum degree at least $2(k-1)$ have a k -block. If the graph G considered is $(k-1)$ -connected, the minimum degree needed comes down to at most $\frac{3}{2}(k-1)$, and further to k if G contains no triangle. In Section 5 we show that graphs of average degree at least $3(k-1)$ contain a k -block. In Section 6 we clarify the relationship between k -blocks and tangles. In Section 7 we present a polynomial-time algorithm that decides whether a given graph has a k -block, and another that finds all the k -blocks in a graph. This latter algorithm gives rise to our duality theorem $\beta = \text{bw}$.

2 Terminology and background

All graph-theoretic terms not defined within this paper are explained in [3]. Given a graph $G = (V, E)$, an ordered pair (A, B) of vertex sets such that $A \cup B = V$ is called a *separation* of G if there is no edge xy with $x \in A \setminus B$ and $y \in B \setminus A$. The sets A, B are the *sides* of this separation. A separation (A, B) such that neither $A \subseteq B$ nor $B \subseteq A$ is a *proper* separation. The *order* of a separation (A, B) is the cardinality of its *separator* $A \cap B$. A separation of order k is called a *k -separation*. A simple calculation yields the following:

Lemma 2.1. *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*

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

- (T1) $V = \bigcup_{t \in T} V_t$;
- (T2) for every edge $e \in G$ there exists a $t \in T$ such that both ends of e lie in V_t ;
- (T3) $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever t_2 lies on the t_1 - t_3 path in T .

¹For example, in a network G one might think of the nodes of a $\beta(G)$ -block as locations to place some particularly important servers that should still be able to communicate with each other when much of the network has failed.

The sets V_t are the *parts* of (T, \mathcal{V}) , its *width* is the number $\max_{t \in T} |V_t| - 1$, 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 (T, \mathcal{V}) (those for which tt' is an edge of T) are its *adhesion sets*; the maximum size of such a set is the *adhesion* of (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 (T3), we have $\mathring{V}_t = V_t \setminus \bigcup_{t' \neq t} V_{t'}$.

Given an edge $e = t_1 t_2$ of T , the two components $T_1 \ni t_1$ and $T_2 \ni t_2$ of $T - e$ define separations (A, B) and (B, A) of G with $A = \bigcup_{t \in T_1} V_t$ and $B = \bigcup_{t \in T_2} V_t$, whose separator is the adhesion set $V_{t_1} \cap V_{t_2}$ [3, Lemma 12.3.1]. We call these the separations *induced* by the tree-decomposition (T, \mathcal{V}) . Note that the adhesion of a tree-decomposition is the maximum of the orders of the separations it induces.

A tree-decomposition *distinguishes* two k -blocks b_1, b_2 if it induces a separation that separates them. It does so *efficiently* if this separation can be chosen of order no larger than the minimum order of a b_1 - b_2 separator in G . The tree-decomposition (T, \mathcal{V}) is *Aut(G)-invariant* if the automorphisms of G act on the set of parts in a way that induces an action on the tree T . The following theorem was proved in [2]:

Theorem 2.2. *For every $k \in \mathbb{N}$, every graph G has an Aut(G)-invariant tree-decomposition of adhesion at most k that efficiently distinguishes all its k -blocks.*

A tree-decomposition (T, \mathcal{V}) of a graph G is *lean* if for any nodes $t_1, t_2 \in T$, 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 $tt' \in t_1 T t_2$ 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 $(< \ell)$ -inseparable for any large ℓ ; see Section 3.)

We call a tree-decomposition (T, \mathcal{V}) *k-lean* if none of its parts contains another, it has adhesion at most k , and for any nodes $t_1, t_2 \in T$, 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 $tt' \in t_1 T t_2$ with $|V_t \cap V_{t'}| < \ell$.

Thomas [15] 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 in [1] to yield the following:

Theorem 2.3. *For every $k \in \mathbb{N}$, every graph has a k -lean tree-decomposition.*

3 Examples of k -blocks

In this section we discuss three different types of k -block.

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

While a k -block as in Example 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 2. Let $k \geq 5$, and let H be a large $(m \times n)$ -grid, with $m, n \geq k^2$ say. Let G be obtained from H by adding a set $X = \{x_1, \dots, x_k\}$ of new vertices, joining each x_i to at least k 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 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 neighbours, because these separate it from all the other vertices. As any k -block has at least $k \geq 5$ 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 independent paths in G , and hence cannot be separated by fewer than k vertices. Hence X is $(< k)$ -inseparable, maximally so, and is thus the only k -block of G .




Figure 1: The six outer vertices form a 6-block

In the discussion of Example 2 we saw that none of the grid vertices lies in a k -block. In particular, the grid itself has no k -block when $k \geq 5$. Since every two inner vertices of the grid, those of degree 4, are joined in the grid by 4 independent paths, they form a (< 4) -inseparable set (which is clearly maximal):

Example 3. The inner vertices of any large grid H form a 4-block in H . However, H has no k -block for any $k \geq 5$.

The k -block defined in Example 2 gives rise to a tangle of large order (see Section 6), 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 6 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 4. Let X be a set of $n \geq k$ isolated vertices. Join every two vertices of X by many (e.g., k) independent paths, making all these internally disjoint. Then X will be a k -block in the resulting graph.

Example 4 differs from Example 2 in that its graph has a tree-decomposition whose only part of order ≥ 3 is X . Unlike the grid in Example 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 5. Given $n \geq k$, consider a tree T in which every non-leaf node has $\binom{n}{k-1}$ 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 $(k-1)$ -subset $S_{t'}$ of V_t by $(k-1)$ 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.

Interestingly, the k -blocks that we shall construct explicitly in our proofs will all be *connected*, i.e., induce connected subgraphs. Thus, our proof techniques seem to be insufficient to detect k -blocks that are disconnected or even independent, such as those in our examples. However, we do not know whether or not this affects the quality of our bounds or just their witnesses:

Problem 1. Does every minimum or average degree bound that forces the existence of a k -block also force the existence of a connected ($< k$)-inseparable set?

Even if the answer to this problem is positive, it will reflect only on how our invariant β relates to the invariants δ and d , and that for some graphs it may be more interesting to relate β to other invariants. The existence of a large k -block in Examples 2 and 4, for instance, will not follow from any theorem relating β to δ or d , since graphs of this type have a bounded average degree independent of k , even in all subgraphs. But they are key examples, which similar results about β and other graph invariants may be able to detect.

4 Minimum degree conditions forcing a k -block

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

Without any further assumptions on G we shall see that $\delta(G) \geq 2(k-1)$ will be enough. If we assume that G is $(k-1)$ -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 $\delta(G) > \frac{3}{2}k - \frac{5}{2}$ suffices. If G is $(k-1)$ -connected but contains no triangle, even $\delta(G) \geq k$ 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 , except in some very special cases that are easy to describe.

Conversely, we construct a $(k-1)$ -connected graph of minimum degree $\lfloor \frac{3}{2}k - \frac{5}{2} \rfloor$ that has no k -block. So our second result above is sharp.

To enhance the readability of both the results and the proofs in this section, we give bounds on δ which force the existence of a $(k+1)$ -block for any $k \geq 0$.

We shall often use the fact that a vertex of G together with k or more of its neighbours forms a $(< k+1)$ -inseparable set as soon as these neighbours are pairwise not separated by k or fewer vertices. Let us state this as a lemma:

Lemma 4.1. *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 $(k+1)$ -block. \square*

Here, then, is our first sufficient condition for the existence of a k -block. It is essentially due to Mader [11, Satz 7'], though with a different proof:

Theorem 4.2. *If $\delta(G) \geq 2k$, then G has a $(k+1)$ -block. This $(k+1)$ -block 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 2.3, G has a k -lean tree-decomposition (T, \mathcal{V}) , say with $\mathcal{V} = (V_t)_{t \in T}$. Pick a leaf t of T . (If 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 (T, \mathcal{V}) , we have $\dot{V}_t = V_t \setminus A_t \neq \emptyset$ by (T3). By our degree assumption and $|A_t| \leq k$, every vertex in \dot{V}_t has k neighbours in \dot{V}_t . Thus, $|\dot{V}_t| \geq k+1 \geq 2$.

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

One might expect that our lower bound for the size of the $(k+1)$ -block B found in the proof of Theorem 4.2 can be increased by proving that B must contain the adhesion set of the part V_t containing it. While we can indeed raise the bound a little (by at least 1, but we do not know how much at most), we show in Section 8 that B can lie entirely in the interior of V_t .

We also do not know whether the degree bound of $\delta(G) \geq 2k$ in Theorem 4.2 is sharp. The largest minimum degree known of a graph without a $(k+1)$ -block is $\lfloor \frac{3}{2}k - 1 \rfloor$. This graph (Example 6 below) is k -connected, and we shall see that k -connected graphs of larger minimum degree do have $(k+1)$ -blocks (Theorem 4.6). Whether or not graphs of minimum degree between $\frac{3}{2}k - 1$ and $2k$ and connectivity $< k$ must have $(k+1)$ -blocks is unknown to us:

Problem 2. Given $k \in \mathbb{N}$, determine the smallest value δ_k of δ such that every graph of minimum degree at least δ has a k -block.

It is also conceivable that the smallest minimum degree that will force a *connected* $(k+1)$ -block – or at least a connected $(< k+1)$ -inseparable set, as found by our proof of Theorem 4.2 – is indeed $2k$ but possibly disconnected $(k+1)$ -blocks can be forced by a smaller value of δ (compare Problem 1).

The degree bound of Theorem 4.2 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. 2) 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 2: The separation (A, B) is T-shaped

Lemma 4.3. *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

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

When a k -separation (A, B) is T-shaped, no $(k + 1)$ -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 $(k + 1)$ -block. Conversely, one may ask whether every proper k -separation (A, B) in a k -connected graph such that A contains no $(k + 1)$ -block must be T-shaped, or at least give rise to a T-shaped k -separation (A', B') with $A' \subseteq A$. This, however, is not true: some counterexamples are given in Section 8.

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 $(k + 1)$ -block. More precisely, we have the following:

Lemma 4.4. *If G is k -connected, the following statements are equivalent:*

- (i) *every proper k -separation of G separates two $(k + 1)$ -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 $(k + 1)$ -block. If A contains a $(k + 1)$ -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 $(k + 1)$ -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 $(k + 1)$ -block, Lemma 4.1 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 2.1 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 $(k + 1)$ -blocks we shall find.

Lemma 4.5. *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 $(k + 1)$ -block or G has two $(k + 1)$ -blocks of size at least $\min\{|A| : (A, B) \text{ is a proper } k\text{-separation}\}$ that are connected in G .*

Proof. By assumption and Lemma 4.3, G has no T-shaped k -separation, so by Lemma 4.4 every side of a proper k -separation contains a $(k+1)$ -block.

By Theorem 2.3, G has a k -lean tree-decomposition (T, \mathcal{V}) , with $\mathcal{V} = (V_t)_{t \in T}$ say. Unless V is a $(k+1)$ -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 T has at least two leaves, and for every leaf t the separation $(A, B) := (V_t, \bigcup_{t' \neq t} V_{t'})$ is a proper k -separation. It thus suffices to show that $A = V_t$ is a $(k+1)$ -block; it will clearly be connected (as in the proof of Theorem 4.2).

As remarked at the start of the proof, there exists a $(k+1)$ -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 (T, \mathcal{V}) for $V_t = A$.

So the order of $(A \cap C, B \cup D)$ is at least $k+1$. By Lemma 2.1, 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 $(k+1)$ -block in $B \cap D$.

The $(k+1)$ -block X' which D contains (see earlier) thus lies in $D \cap A$. So A contains two $(k+1)$ -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 (T, \mathcal{V}) . \square

Theorem 4.6. *If G is k -connected and $\delta(G) > \frac{3}{2}k - 1$, then either V is a $(k+1)$ -block or G has at least two $(k+1)$ -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 4.5. \square

To show that the degree bound in Theorem 4.6 is sharp, let us construct a k -connected graph H with $\delta(H) = \lfloor \frac{3}{2}k - 1 \rfloor$ that has no $(k+1)$ -block.

Example 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

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

see Figure 3. This graph H is k -connected and has minimum degree $\frac{3}{2}k - 1$. But it contains no $(k+1)$ -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' .




Figure 3: A k -connected graph without a $(k+1)$ -block

Clearly, H is again k -connected and has minimum degree $\lfloor \frac{3}{2}k - 1 \rfloor$, and it has no $(k+1)$ -block since H has no k -block.

Our next example shows that the connectivity bound in Theorem 4.6 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 $(k+1)$ -blocks have size $k+1 < \delta(H) + 1$.

Example 7. Let H_n be as in Example 6. Let H be obtained from H_n by replacing the degree-two vertices of H_n by complete graphs of order $(k+1)/2$ 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 $(k+1)$ -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 4.6 is necessary if we just want to force any $(k+1)$ -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 $(k+1)$ -block comes down to $k+1$, and the $(k+1)$ -blocks we find are also larger:

Theorem 4.7. *If G is k -connected, $\delta(G) \geq k+1$, and G contains no triangle, then either V is a $(k+1)$ -block or G has at least two $(k+1)$ -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 4.5 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 4.7 is sharp, because of the following observation:

Proposition 4.8. *If G is k -connected and k -regular, then G has no $(k+1)$ -block unless $G = K^{k+1}$ (which contains a triangle).*

Proof. Suppose G has a $(k+1)$ -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 $(k+1)$ -blocks, regardless of its degree:

Theorem 4.9. *If $\kappa(G) = k \geq 1$ and G is transitive, then G has no $(k+1)$ -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 $(k+1)$ -block. Let us show that G has no $(k+1)$ -block at all.

If G has a $(k+1)$ -block, it has at least two, since V is not a $(k+1)$ -block but every vertex lies in a $(k+1)$ -block, by transitivity. Hence any tree-decomposition that distinguishes all the $(k+1)$ -blocks of G has at least two parts. By Theorem 2.2 there exists such a tree-decomposition (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 T and t' is its neighbour in T , then every vertex in $V_t \setminus V_{t'}$ lies in no other part than V_t (see Section 2). Hence no automorphism of G maps u to v , a contradiction to the transitivity of G . \square

Theorems 4.6 and 4.9 together imply a well-known theorem of Mader [9] and Watkins [16], which says that every transitive graph of connectivity k has minimum degree at most $\frac{3}{2}k - 1$.

5 Average degree conditions forcing a k -block

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 in the previous section, we shall assume that $k \geq 0$ and consider $(k+1)$ -blocks, to improve readability.

As remarked in the introduction, Mader [10] proved that if $d(G) \geq 4k$ then G has a $(k+1)$ -connected subgraph. The vertex set of such a subgraph is $(< k+1)$ -inseparable, and hence extends to a $(k+1)$ -block of G . Our first aim will be to show that if we seek to force a $(k+1)$ -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 5.1. *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 5.2. *If $d(G) \geq 3k$, then G has a $(k + 1)$ -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 5.1, 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 (T, \mathcal{V}) be a k -lean tree-decomposition of G , with $\mathcal{V} = (V_t)_{t \in T}$. Pick a leaf t of T . (If T has only one node, let t be this node.) Then $V_t \neq \emptyset$ by (T3), since V_t is not contained in any other part of \mathcal{V} .

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

$$|E(\hat{V}_t, V)| \leq \frac{1}{2}|\hat{V}_t|^2 + k|\hat{V}_t| \leq |\hat{V}_t|(|\hat{V}_t|/2 + k) \leq \frac{3}{2}k|\hat{V}_t|,$$

which contradicts Lemma 5.1. So $|\hat{V}_t| \geq k + 1 \geq 2$. The set \hat{V}_t extends to a $(k + 1)$ -block $B \subseteq V_t$ with the desired properties as in the proof of Theorem 4.2. \square

Since our graph of Example 6 contains no $(k + 1)$ -block, its average degree is a strict lower bound for the minimum average degree that forces a $(k + 1)$ -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 $(k + 1)$ -block thus lies somewhere between $2k - 1$ and $3k$.

Problem 3. Given $k \in \mathbb{N}$, determine the smallest value d_k of d such that every graph of average degree at least d has a k -block.

As we have seen, an average degree of $3k$ is sufficient to force a graph to contain a $(k + 1)$ -block. If we ask only that the graph should have a minor that contains a $(k + 1)$ -block, then a smaller average degree suffices:

Theorem 5.3. *If $d(G) \geq 2(k - 1) > 0$, then G has a minor containing a $(k + 1)$ -block. This $(k + 1)$ -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 5.1.

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 (T, \mathcal{V}) be a k -lean tree-decomposition of G , with $\mathcal{V} = (V_t)_{t \in T}$ say, and let t be a leaf of T . (If T has only one node, let t be this node.) We shall prove that V_t is $(k+1)$ -inseparable, and hence a $(k+1)$ -block, in G .

As (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 V_t , as otherwise $X := A_t \setminus \{a\}$ would separate A_t from every set $X \cup \{v\}$ with $v \in V_t$, 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 $V_t \neq \emptyset$ 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 6 have average degrees of at least $2k-1-\epsilon$. So these graphs show that obtaining a $(k+1)$ -block in G is indeed harder than obtaining a $(k+1)$ -block in a minor of G , which these graphs must have by Theorem 5.3. (And they do: they even have $K^{3k/2}$ -minors.)

6 Blocks and tangles

In this section we compare k -blocks with tangles, as introduced by Robertson and Seymour [14]. Our reason for doing so is that both notions have been advanced as possible approximations to the elusive “ $(k+1)$ -connected pieces” into which one might wish to decompose a k -connected graph, in analogy to its tree-like block-cutvertex decomposition (for $k=1$), or to Tutte’s tree-decomposition of 2-connected graphs into 3-connected torsos (for $k=2$) [13, 2].

Let us say that a set θ of separations of order at most k of a graph $G = (V, E)$ is a *tangle of order k* of G if

- ($\theta 1$) for every separation (A, B) of order $< k$ of G either (A, B) or (B, A) is in θ ;
- ($\theta 2$) for all $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \theta$ we have $G[A_1] \cup G[A_2] \cup G[A_3] \neq G$.

It is straightforward to verify that this notion of a tangle is consistent with the one given in [14].

Given a tangle θ , we think of the side A of a separation $(A, B) \in \theta$ as the *small side* of (A, B) , and of B as its *large side*. (Thus, axiom ($\theta 2$) says that G is not the union of the subgraphs induced by at most three small sides.) If a set X of vertices lies in the large side of every separation in θ but not in the small side, we say that X *gives rise to* or *defines* the tangle θ .

If X is a (k) -inseparable set of vertices, it clearly lies in exactly one of the two sides of any separation of order $< k$. Hence if we define θ as the set of those separations (A, B) of order $< k$ for which $X \subseteq B$, then θ satisfies ($\theta 1$), and V is not a union of at most two small sides of separations in θ . But it might be the

union of three small sides, and indeed θ may fail to satisfy $(\theta 2)$. So X might, or might not, define a tangle of order at most k .

An $(n \times n)$ -grid minor of G , with $n \geq k$, also gives rise to a tangle of order k in G , but in a weaker sense: for every separation (A, B) of G of order less than k , exactly one side meets a branch set of every *cross* of the grid, a union of one column and one row. (Indeed, since crosses are connected and every two crosses meet, we cannot have one cross in $A \setminus B$ and another in $B \setminus A$.)

Since G can contain a large grid without containing a k -block (Example 3), it can thus have a large-order tangle but fail to have a k -block for any $k \geq 5$. Conversely, Examples 4 and 5 show that G can have k -blocks for arbitrarily large k without containing any tangle (other than those of order $\leq \kappa(G)$, in which the large side of every separation is all of V). For example, if G is a subdivided K^n with $n \geq k+1$, then its branch vertices form a k -block X , but when $n \leq \frac{3}{2}(k-1)$ the separations of order $< k$ whose large sides contain X do not form a tangle, since G is the union of three small sides of such separations (each with a separator consisting of two thirds of the branch vertices; compare [14, (4.4)]).

Any k -block of size $> \frac{3}{2}(k-1)$, however, does give rise to a tangle of order k :

Theorem 6.1. *Every $(< k)$ -inseparable set of more than $\frac{3}{2}(k-1)$ vertices in $G = (V, E)$ defines a tangle of order k .*

Proof. Let X be a $(< k)$ -inseparable set of more than $\frac{3}{2}(k-1)$ vertices, and consider the set θ of all separations (A, B) of order less than k with $X \subseteq B$. We show that θ is a tangle. As no two vertices of X can be separated by a separation in θ , it satisfies $(\theta 1)$. For a proof of $(\theta 2)$, it suffices to consider three arbitrary separations $(A_1, B_1), (A_2, B_2), (A_3, B_3)$ in θ and show that

$$E(A_1) \cup E(A_2) \cup E(A_3) \not\supseteq E, \quad (*)$$

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

As $|X| > \frac{3}{2}(k-1)$, there is a vertex $v \in X$ 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 $X \subseteq B_1$, there is another vertex $w \neq v$ in $X \setminus A_1$. As v and w lie in X , 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

7 Finding k -blocks in polynomial time

We consider graphs $G = (V, E)$, with n vertices and m edges, say, and positive integers $k < n$. We shall present a simple algorithm that finds all the k -blocks of G in time polynomial in n, m and k . We start our algorithm with the following step, which we call *pre-processing*.

For two vertices x, y of G let $\kappa(x, y)$ denote the smallest size of a set of other vertices that separates x from y in G . We construct a graph H_k from G by

adding, for every pair of non-adjacent vertices x, y , the edge xy if $\kappa(x, y) \geq k$, that is, if x and y cannot be separated by fewer than k other vertices. Moreover, we label every non-edge xy of H_k by some separation of order $\kappa(x, y) < k$ that separates x from y in G . This completes the pre-processing.

Note that all separations of order $< k$ of G are still separations of H_k , and that the k -blocks of G are the vertex sets of the maximal cliques of order $\geq k$ in H_k .

Lemma 7.1. *The pre-processing has running time $O(\min\{k, \sqrt{n}\} \cdot m \cdot n^2)$.*

Proof. We turn the problem of finding a minimal vertex separator between two vertices into one of finding a minimal edge cut between them. This is done in the usual way (see e.g. Even [6]) by constructing a unit-capacity network G' from G with $n' = 2\tilde{n}$ vertices and $m' = 2m + \tilde{n}$ directed edges, where $\tilde{n} = O(m)$ is the number of non-isolated vertices of G .

For every non-edge xy of G we start Dinitz's algorithm (DA) on G' , which is designed to find an x - y separation of order $\kappa(x, y)$. If DA completes k iterations of its 'inner loop' (finding an augmenting path), then $\kappa(x, y) \geq k$; we then stop DA and let xy be an edge of H_k . Otherwise DA returns a separation (A, B) of order $< k$; we then keep xy as a non-edge of H_k and label it by (A, B) . Since the inner loop has time complexity $O(m') = O(m)$ and DA has an overall time complexity of $O(\sqrt{n'} \cdot m') = O(\sqrt{n} \cdot m)$ (see e.g. [8]), this establishes the desired bound. \square

Now we describe the *main part of the algorithm*. We shall construct a rooted tree T , inductively by adding children to leaves of the tree constructed so far. We maintain two lists: a list \mathcal{L} of some of the leaves of the current tree, and a list \mathcal{B} of subsets of V . We shall change \mathcal{L} by either deleting its last element or replacing it with two new elements that will be its children in our tree. Whenever we add an element t to \mathcal{L} in this way, we assign it a set $X_t \subseteq V$. Think of the current list \mathcal{L} as containing those t whose X_t we still plan to scan for k -blocks of G , and of \mathcal{B} as the set of k -blocks found so far.

We start with a singleton list $\mathcal{L} = (r)$ and $\mathcal{B} = \emptyset$, putting $X_r = V$.

At a given step, stop with output \mathcal{B} if \mathcal{L} is empty; otherwise consider the last element t of \mathcal{L} . If $|X_t| < k$, delete t from \mathcal{L} and do nothing further at this step.

Assume now that $|X_t| \geq k$. If X_t induces a complete subgraph in H_k , add X_t to \mathcal{B} , delete t from \mathcal{L} , and do nothing further at this step.

If not, find vertices $x, y \in X_t$ that are not adjacent in H_k . At pre-processing, we labeled the non-edge xy with a separation (A, B) of order $< k$ that separates x from y in G (and in H_k). Replace t in \mathcal{L} by two new elements t' and t'' , making them children of t in the tree under construction, and let $X_{t'} = X_t \cap A$ and $X_{t''} = X_t \cap B$. If $|X_t| > k$, do nothing further at this step. If $|X_t| = k$, then both $X_{t'}$ and $X_{t''}$ have size $< k$; we delete t' and t'' again from \mathcal{L} and do nothing further in this step.

This completes the description of the main part of the algorithm. Let T be the tree with root r that the algorithm constructed: its nodes are those t that were in \mathcal{L} at some point, and its edges were defined as nodes were added to \mathcal{L} .

Proposition 7.2. *The main part of the algorithm stops with output \mathcal{B} the set of k -blocks of G .*

Proof. The algorithm clearly stops with \mathcal{B} the set of vertex sets of the maximal cliques of H_k that have order $\geq k$. These are the k -blocks of G , by definition of H_k . \square

To analyse running time, we shall need a lemma that is easily proved by induction. A *leaf* in a rooted tree is a node that has no children, and a *branching node* is one that has at least two children.

Lemma 7.3. *Every rooted tree has more leaves than branching nodes.* \square

Lemma 7.4. *The main part of the algorithm stops after at most $4(n-k)$ steps. Its total running time is $O(\min\{m, n\} \cdot n^2)$.*

Proof. Each step takes $O(n^2)$ time, the main task being to check whether $H_k[X_t]$ is complete. It thus suffices to show that there are no more than $4(n-k)$ steps as long as $n \leq 2m$, which can be achieved by deleting isolated vertices.

At every step except the last (when $\mathcal{L} = \emptyset$) we considered the last element t of \mathcal{L} , which was subsequently deleted or replaced and thus never considered again. Every such t is a node of the tree T' obtained from T by deleting the children of nodes t with $|X_t| = k$. (Recall that such children t', t'' were deleted again immediately after they were created, so they do not give rise to a step of the algorithm.) Our aim, therefore, is to show that $|T'| \leq 4(n-k) - 1$.

By Lemma 7.3 it suffices to show that T' has at most $2(n-k)$ leaves. As $n \geq k+1$, this is the case if T' consists only of its root r . If not, then r is a branching node of T' . It thus suffices to show that below every branching node t of T' there are at most $2(|X_t| - k)$ leaves; for $t = r$ this will yield the desired result.

By definition of T' , branching nodes t of T' satisfy $|X_t| \geq k+1$. So our assertion holds if the two children of t are leaves. Assuming inductively that the children t' and t'' of t satisfy the assertion (unless they are leaves), we find that, with $X_{t'} = X_t \cap A$ and $X_{t''} = X_t \cap B$ for some ($< k$)-separation (A, B) of G as in the description of the algorithm, the number of leaves below t is at most

$$2(|X_t \cap A| - k) + 2(|X_t \cap B| - k) \leq 2(|X_t| + (k-1) - 2k) \leq 2(|X_t| - k)$$

if neither t' nor t'' is a leaf, and at most

$$1 + 2(|X_t \cap B| - k) \leq 2(|X_t| - k)$$

if t' is a leaf but t'' is not (say), since $X_t \setminus B \neq \emptyset$ by the choice of (A, B) . \square

Putting Lemmas 7.1 and 7.4 together, we obtain the following:

Theorem 7.5. *There is an $O(\min\{k, \sqrt{n}\} \cdot m \cdot n^2)$ -time algorithm that finds, for any graph G with n vertices and m edges and any fixed $k < n$, all the k -blocks in G .* \square

Our algorithm can easily be adapted to find the k -blocks of G for all values of k at once. To do this, we run our pre-processing just once to construct the graph H_n , all whose non-edges xy are labeled by an x - y separation of minimum order and its value $\kappa(x, y)$. We can then use this information at the start of the proof of Lemma 7.4, when we check whether $H_k[X_t]$ is complete, leaving the running time of the main part of the algorithm at $O(n^3)$ as in Lemma 7.4. Running it separately once for each $k < n$, we obtain with Lemma 7.1:

Theorem 7.6. *There is an $O(\max\{m\sqrt{n}n^2, n^4\})$ algorithm that finds, for any graph G with n vertices and m edges, all the k -blocks of G (for all k). \square*

Perhaps this running time can be improved if the trees T_k exhibiting the k -blocks are constructed simultaneously, e.g. by using separations of order ℓ for all T_k with $\ell < k$.

The mere decision problem of whether G has a k -block does not need our pre-processing, which makes the algorithm faster:

Theorem 7.7. *For fixed k , deciding whether a graph with n vertices and m edges has a k -block has time complexity $O(mn + n^2)$.*

Proof. Given k and a graph G , we shall find either a ($< k$)-inseparable set of vertices in G (which we know extends to a k -block) or a set S of at most $2(n - k) - 1$ separations of order $< k$ such that among any k vertices in G some two are separated by a separation in S (in which case G has no k -block).

Starting with $X = V(G)$, we pick a k -set of vertices in X and test whether any two vertices in this set are separated by a ($< k$)-separation (A, B) in G . If not, we have found a ($< k$)-inseparable set of vertices and stop with a yes-answer. Otherwise we iterate with $X = A$ and $X = B$.

Every separation found by the algorithm corresponds to a branching node of T . All these are nodes of T' , of which there are at most $4(n - k) - 1$ (see the proof of Lemma 7.4). Testing whether a given pair of vertices is separated by some ($< k$)-separation of G takes at most k runs of the inner loop of Dinitz's algorithm (which takes $O(m + n)$ time), and we test at most $\binom{k}{2}$ pairs of vertices in X . \square

Let us say that a set S of ($< k$)-separations in G *witnesses* that G has no k -block if among every k vertices of G some two are separated by a separation in S . Trivially, if G has no k -block then this is witnessed by some $O(n^2)$ separations. The proof of Theorem 7.7 shows that this bound can be made linear:

Corollary 7.8. *Whenever a graph of order n has no k -block, there is a set of at most $4(n - k) - 1$ separations witnessing this. \square*

Let us call any tree T as in our main algorithm (at any stage), with each of its branching nodes t labelled by a separation $(A, B)_t$ of G that separates some two vertices of X_t , a *block-decomposition* of G . The sets X_t with t a leaf will be called its *leaf sets*.

The *adhesion* of a block-decomposition is the maximum order of the separations $(A, B)_t$. A block-decomposition is *k-complete* if it has adhesion $< k$ and every leaf set is $(< k)$ -inseparable or has size $< k$. The *width* of a block-decomposition is the maximum order of a leaf set X_t . The *block-width* $\text{bw}(G)$ of G is the least k such that G has a block-decomposition of adhesion and width both at most k .

Having block-width $< k$ can be viewed as dual to containing a k -block, much as having tree-width $< k - 1$ is dual to containing a haven or bramble of order k , and having branch-width $< k$ is dual to containing a tangle of order k . Indeed, we have shown the following:

Theorem 7.9. *Let $\mathcal{D} = (T; (A, B)_t, t \in T)$ be a block-decomposition of a graph G , and let $k \in \mathbb{N}$.*

- (i) *Every edge of G has both ends in some leaf set of T .*
- (ii) *If \mathcal{D} has adhesion $< k$, then any k -block of G is contained in a leaf set of T .*
- (iii) *If \mathcal{D} is k -complete, then every k -block of G is a leaf set, and all other leaf sets have size $< k$. □*

Theorem 7.9 implies that G has a block-decomposition of adhesion and width both at most k if and only if G has no $(k + 1)$ -block. The least such k clearly equals the greatest k such that G has a k -block, its *block number* $\beta(G)$:

Corollary 7.10. *Every finite graph G satisfies $\beta(G) = \text{bw}(G)$. □*

By Theorem 7.7 and its proof, we obtain the following complexity bound:

Corollary 7.11. *Deciding whether a graph with n vertices and m edges has block-width $< k$, for k fixed, has time complexity $O(mn + n^2)$. □*

For k variable, the proof of Theorem 7.7 yields a complexity of $O(k^3(m + n)(n - k))$. Alternatively, we can use pre-processing to obtain $O(\min\{k, \sqrt{n}\} \cdot m \cdot n^2)$ by Theorem 7.5.

The above duality between the block number and the block-width of a graph is formally reminiscent of the various known dualities for other width parameters, such as the tree-width, branch-width, path-width, rank-width, carving-width or clique-width of a graph. The ‘width’ to which these parameters refer, however, is usually that of a tree-like decomposition of the graph itself, which exhibits that it structurally resembles that tree. In our block-decompositions, on the other hand, the tree T merely indicates a recursion by which the graph can be decomposed into small sets: the separations used to achieve this, though of small order, will not in general be nested, and the structure of G will not in any intuitive sense be similar to that of T .

In [5], Diestel and Oum give a structural duality theorem for k -blocks in the sense of those traditional width parameters. The graph structure that is shown to witness the absence of a k -block is not a tree-structure, but one modelled on more general (though still tree-like) graphs. Whether or not a structural duality between k -blocks and tree-like decompositions exists remains an open

problem. It has been formalized, and stated explicitly [5], with reference to a fundamental structural duality theorem between tangle-like ‘dense objects’ and tree-like decompositions, which implies all the traditional duality theorems for width parameters [4] but does not yield a duality theorem for k -blocks.

8 Further examples

In this section we discuss several examples dealing with certain situations of our results. In particular, we will describe one example that shows that the $(k + 1)$ -block found in Theorem 4.2 need not contain any vertex of the adhesion set that lies in the same part of the tree-decomposition, and we will describe two examples dealing with the notion of T-shaped and Lemma 4.4. All these examples are included only in this extended version of this paper.

Recall that in the proof of Theorem 4.2 we considered a k -lean tree tree-decomposition (T, \mathcal{V}) of a graph G with $\delta(G) \geq 2k$ and showed for each leaf t of T that V_t includes a $(k + 1)$ -block b . We now give an example where the adhesion set $V_t \cap V_{t'}$ lies completely outside b , where t' is the neighbour of t in T .




Figure 4: S lies outside the 4-block containing the K^7

Example 8. Let G be the graph in Figure 4 and let (T, \mathcal{V}) be the tree-decomposition with adhesion sets S and those 2-separators that contain one vertex in S and the lowest vertex. So T is a star with 4 leaves. It is not hard to show that (T, \mathcal{V}) is 3-lean. For every vertex x of the adhesion set S inside the upper part V_t , its two neighbours in V_t together with the bottom vertex separate it from any vertex in $\overset{\circ}{V}_t$ but its neighbours. Hence x does not lie in the 4-block b that contains $\overset{\circ}{V}_t$. As no vertex of S lies in b , we conclude $\overset{\circ}{V}_t = b$.

Our next example shows that a local version of Lemma 4.4 as discussed just before the lemma is false. We considered there the question of whether every proper k -separation (A, B) in a k -connected graph such that A contains no $(k + 1)$ -block must be T-shaped, at least if A is minimal as above.

Example 9. Let $k = 6$, and let G be the complement of the disjoint union of three induced paths P_1, P_2, P_3 of length 2. Then each of three sets $V(P_i)$

is separated by the union of the other two. Hence any 7-block misses a vertex from each P_i and thus has at most 6 vertices. Hence, G has no 7-block.

But G is 6-connected, and its only proper 6-separations (A, B) have the form that either $A \setminus B$ consists of the ends of some P_i and $B \setminus A$ of its inner vertex, or vice versa. Let (A, B) be a 6-separation of the first kind. Obviously, A is minimal such that (A, B) , for some B , is a proper 6-separation.

To show that (A, B) is not T-shaped, suppose it is, and let this be witnessed by another proper 6-separation (C, D) . Then (C, D) is neither (A, B) nor (B, A) . So the separators $A \cap B$ and $C \cap D$ meet in exactly one $V(P_i)$, say in $V(P_1)$. Then $C \cap D$ contains $V(P_2)$, say, while $A \cap B$ contains $V(P_3)$. By assumption, the ends of P_2 lie in $A \setminus B$. If the ends of P_3 lie in $C \setminus D$, say, we have $|A \cap C| = 7$. This contradicts the choice of (C, D) , so (A, B) is not T-shaped.

So our envisaged local version of Lemma 4.4 is false. Since $|A| = 8 \leq \frac{3}{2}k$ in the above example, we could not simply use Lemma 4.3 to show that (A, B) is not T-shaped. In our next example A is larger, so that we can.

Example 10. Let G be the graph of Figure 5. It is 5-connected but has no 6-block. Let A be the vertex set that consists of the vertices of the upper three K^5 s, and let B be the union of the vertex sets of the lower three complete graphs. Then (A, B) is a proper 5-separation, with A minimal. By Lemma 4.3, (A, B) is not T-shaped.

By Lemma 4.4, however, both these examples must have some T-shaped k -separation. In Example 9, the separation (B, A) is T-shaped. In Example 10, the separation (A', B') where A' consists of the two leftmost complete graphs and B' of the other four, is T-shaped.




Figure 5: A 5-connected graph without a 6-block

9 Acknowledgements

We thank Jens Schmidt for pointing out reference [11] for Theorem 4.2, Matthias Kriesell for advice on the connectivity of triangle-free graphs, Paul Seymour for

suggesting an algorithm of the kind indicated in Section 7, and Sang-il Oum for pointing out Corollary 7.8 and Theorem 7.9.

References

- [1] P. Bellenbaum and R. Diestel. Two short proofs concerning tree-decompositions. *Comb., Probab. Comput.*, 11:1–7, 2002.
- [2] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein. Connectivity and tree structure in finite graphs. *Combinatorica*, 34(1):1–35, 2014.
- [3] R. Diestel. *Graph Theory*. Springer, 4th edition, 2010.
- [4] R. Diestel and S. Oum. Unifying duality theorems for width parameters, I. Weak and strong duality. arXiv:1406.3797, 2014.
- [5] R. Diestel and S. Oum. Unifying duality theorems for width parameters, II. General duality. arXiv:1406.3798, 2014.
- [6] S. Even. *Graph Algorithms*. Computer Science Press, 1979.
- [7] F. Hundertmark. Profiles. An algebraic approach to combinatorial connectivity. arXiv:1110.6207, 2011.
- [8] F. Kammer and H. Täubig. Graph connectivity. www.informatik.uni-augsburg.de/thi/personen/kammer/Graph_Connectivity.pdf, 2004.
- [9] W. Mader. Über den Zusammenhang symmetrischer Graphen. *Arch. Math. (Basel)*, 21:331–336, 1970.
- [10] W. Mader. Existenz n -fach zusammenhängender Teilgraphen in Graphen mit genügend hoher Kantendichte. *Abh. Math. Sem. Univ. Hamburg*, 37:86–97, 1972.
- [11] W. Mader. Kreuzungsfreie a, b -wege in endlichen graphen. *Abh. Math. Sem. Univ. Hamburg*, 42:187–204, 1974.
- [12] W. Mader. Über n -fach zusammenhängende Eckenmengen in Graphen. *J. Combin. Theory (Series B)*, 25:74–93, 1978.
- [13] B.A. Reed. Tree width and tangles: a new connectivity measure and some applications. In R.A. Bailey, editor, *Surveys in Combinatorics*. Cambridge Univ. Press, 1997.
- [14] N. Robertson and P.D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)*, 52:153–190, 1991.
- [15] R. Thomas. A Menger-like property of tree-width; the finite case. *J. Combin. Theory (Series B)*, 48:67–76, 1990.
- [16] M.E. Watkins. Connectivity of transitive graphs. *J. Combin. Theory*, 8:23–29, 1970.