Two short proofs concerning tree-decompositions

Patrick Bellenbaum & Reinhard Diestel

We give short proofs of the following two results: of Thomas's theorem that every finite graph has a linked tree-decomposition of width no greater than its tree-width, and of the 'tree-width duality theorem' of Seymour and Thomas that the tree-width of a finite graph is exactly one less than the largest order of its brambles.

1. Introduction

The main purpose of this note is to give a short proof of Thomas's theorem that every finite graph has a linked tree-decomposition of width no greater than its tree-width [9]. This is a useful tool in the theory of tree-decompositions; for example, it is a key lemma in Robertson & Seymour's proof of the fact that every set of graphs of bounded tree-width is well-quasi-ordered [7]. This latter result is the starting point for the proof of their graph minor theorem, see [2]. It is also the first step in the now available short proof of the 'general Kuratowski theorem' that embeddability in any fixed surface is characterized by finitely many forbidden minors (combine it with [3] and either [5] or [10]), a main corollary of the graph minor theorem.

Another (more constructive) short proof of Thomas's theorem has been given in [1]. An analogous result for branch-width was obtained by Geelen, Gerards & Whittle [4], also with a short and simple proof. The 'branch-width' of a graph is a parameter closely related to tree-width but not 1–1 translatable, so the result proved in [4] does not imply Thomas's theorem as reproved in this paper. But [4] does give a complete short proof (including the WQO part) of the above-mentioned result from [7], where exact bounds for Thomas's theorem are not required.

Our proof of Thomas's theorem differs from the original in that we use a simpler induction parameter. This simplifies the induction step: we have less to verify, and the presentation becomes considerably less technical. The tree-decomposition we use, however, is the same as in [9]. Essentially, it is constructed recursively from tree-decompositions of two subgraphs by adding their intersection as a new part to both decompositions, to serve as the common interface required for their amalgamation. Whether or not this can be done without increasing the width is a key issue in the study of tree-decompositions more generally. Thomas's sufficient condition under which this technique can be applied is a contribution of independent use and interest, and so we have extracted it into a separate lemma (Lemma 2).

This lemma (with some inessential additional details) has been used again before, in Seymour & Thomas's proof of what Reed [6] has called their *tree-width duality* theorem [8]. A streamlined proof of this result has already ap-

peared in $[2]$, with the lemma incorporated. Since we need the lemma by itself in our different context here, we take the opportunity to include the short derivation of the duality theorem from it as well (Theorem 5).

All the graphs we consider are finite. Unless otherwise specified our terminology follows $[2]$, and we assume familiarity with the basic theory of tree-decompositions as covered there. In particular, we shall freely use the following separation lemma:

Lemma 1. *Let $(V_t)_{t \in T}$ be a tree-decomposition of a graph G, let $t_1 t_2$ be an edge of T, and let T_1 and T_2 be the components of $T - t_1 t_2$. Then $V_{t_1} \cap V_{t_2}$ separates $\bigcup_{t \in T_1} V_t$ from $\bigcup_{t \in T_2} V_t$ in G.* □

2. Amalgamating tree-decompositions

Let G be a graph and $X \subseteq V(G)$ a separating set of vertices. A basic technique in the study of tree-decompositions is to try to amalgamate given tree-decompositions of the subgraphs $H = G[C \cup X]$, where C ranges over the components of $G - X$, into an overall tree-decomposition of G. This is straightforward if X is contained in a part in each of those decompositions. If X is not contained in a part of the given decomposition of H, we can alter that decomposition and force X to become included in a part; to maintain the decomposition axiom (T3), however, this will typically involve the inclusion of the vertices from X in some of the other parts as well, which can increase the width of that decomposition.

The lemma we prove in this section offers a sufficient condition for when X can be incorporated into a part of a given decomposition of H without increasing its width. The decomposition of H will be given as induced by another tree-decomposition $\mathcal{D}$ of all of G, and the sufficient condition will involve the position of X within $\mathcal{D}$.

Slightly more generally let C be a union of components of $G - X$ and $H := G[C \cup X]$. Let $\mathcal{D} = (V_t)_{t \in T}$ be a tree-decomposition of G, and $s \in T$. For each $x \in X$ pick a 'home' node $t_x \in T$ with $x \in V_{t_x}$, and for all $t \in T$ put

$$W_t := \big(V_t \cap V(H)\big) \cup \{\, x \in X \mid t \in t_x T s \,\}.$$

Clearly $X \subseteq W_s$, and it is easy to check that $\mathcal{D}_s(H) := (W_t)_{t \in T}$ is a tree-decomposition of H. (In fact, it is the tree-decomposition obtained from the decomposition which $\mathcal{D}$ induces on H by forcing $X \subseteq W_s$ and repairing (T3) with the minimum addition of vertices to existing parts.)

Lemma 2. *If $G - C$ contains a set $\{\, P_x \mid x \in X \,\}$ of disjoint X–V_s paths with $x \in P_x$ for all x, then $|W_t| \leqslant |V_t|$ for all $t \in T$.*

Proof. For each $x \in X$ with $x \in W_t \smallsetminus V_t$ we have $t \in t_x T s$, so V_t separates V_{t_x} from V_s by Lemma 1 and hence contains some other vertex of P_x. That vertex does not lie in W_t, because $W_t \smallsetminus X \subseteq V(C)$ while $P_x \subseteq G - C$. □

2

3. Linked and lean tree-decompositions

A tree-decomposition $\mathcal{D} = (V_t)_{t \in T}$ is called *linked* if, in addition to the usual axioms (T1)–(T3), it satisfies the following:

(T4) Given any $k \in \mathbb{N}$ and $t_1, t_2 \in T$, either G contains k disjoint V_{t_1}–V_{t_2} paths or there exists a $t \in t_1 T t_2$ such that $|V_t| < k$.

Let us call the decomposition $\mathcal{D}$ *lean* if, in addition to (T1)–(T3), it satisfies

(T4′) Given $t_1, t_2 \in T$ and vertex sets $Z_1 \subseteq V_{t_1}$ and $Z_2 \subseteq V_{t_2}$ such that $|Z_1| = |Z_2| =: k$, either G contains k disjoint Z_1–Z_2 paths or there exists an edge $tt' \in t_1 T t_2$ with $V_t \cap V_{t'} < k$.

It is easy to make a lean tree-decomposition linked without increasing its width: just subdivide every edge $tt' \in T$ by a new node s and add $V_s := V_t \cap V_{t'}$ as a new part. The greater freedom in choosing the sets Z_1 and Z_2 in (T4′), however, makes (T4′) considerably stronger than (T4), especially when $t_1 = t_2$. (For example, the trivial decomposition into one part is always linked, but not lean unless G is complete.) More examples illustrating the difference between linked and lean decompositions, including a justification of term 'lean' based on this difference, can be found in [1].

Theorem 3. *Every graph G has a lean tree-decomposition of width* $\mathrm{tw}(G)$, *the tree-width of G.*

Proof. Let $n := |G|$. Let the *fatness* of a tree-decomposition of G be the n-tuple $(a_0, \ldots, a_{n-1})$, where a_h denotes the number of parts that have exactly $n - h$ vertices. Let $\mathcal{D} = (V_t)_{t \in T}$ be a tree-decomposition of lexicographically minimal fatness. Clearly $\mathcal{D}$ has width $\mathrm{tw}(G)$; we shall prove that $\mathcal{D}$ is lean.

Suppose not. Then there exists a quadruple (t_1, t_2, Z_1, Z_2) as in (T4′) that violates (T4′); we choose one for which t_1 and t_2 have minimum distance in T. Among all the Z_1–Z_2 separators of minimum order in G let X be one that lies 'closest to $t_1 T t_2$' in the sense that $\sum_{x \in X} d_x$ is minimum, where

$$d_x := \min \{ d_T(t, t_x) \mid t \in t_1 T t_2 \text{ and } x \in V_{t_x} \}.$$

Let C_1 denote the union of those components of $G - X$ that meet Z_1, let C_2 be the union of all the other components of $G - X$, and put $H_i := G[C_i \cup X]$ for $i = 1, 2$. By Menger's theorem, there exists a set $\{ P_x \mid x \in X \}$ of disjoint Z_1–Z_2 paths in G such that $x \in P_x$ for all x. Each P_x is the union of two paths $P_x^i \subseteq G - C_i$ meeting exactly in x; the paths P_x^i will be used below to apply Lemma 2.

For every $x \in X$ choose t_x at minimum distance from $t_1 T t_2$ in T so that $x \in V_{t_x}$. Let T^1, T^2 be disjoint copies of T. For $i = 1, 2$ and $t \in T$ let t^i denote the copy of t in T^i, and put $T' := T^1 \cup T^2 + t_2^1 t_1^2$. Let $\mathcal{D}_{t_{3-i}}(H_i) =: (W_t^i)_{t \in T}$ be the tree-decomposition of H_i obtained from $\mathcal{D}$ for $s := t_{3-i}$ as before Lemma 2.

3

Rewriting T as T^i and W_t^i as W_{t^i} we may combine $\mathcal{D}_{t_2}(H_1)$ and $\mathcal{D}_{t_1}(H_2)$ to a tree-decomposition $\mathcal{D}' = (W_t)_{t \in T'}$ of G. (Note that $\mathcal{D}'$ satisfies (T3), because $V(H_1 \cap H_2) = X \subseteq W_{t_1^2} \cap W_{t_2^1}$.)

To complete the proof we show that $\mathcal{D}'$ has smaller fatness than $\mathcal{D}$. To do so, we prove that

$$(\forall t \in T)(\forall i = 1, 2) \left(|W_t^i| = |V_t| \Rightarrow W_t^{3-i} \subseteq X \right) \tag{1}$$

and

$$(\exists t \in t_1 T t_2) \left(|W_t^1|, |W_t^2| < |V_t| \right). \tag{2}$$

To see why this suffices, note that $|W_t^i| \leqslant |V_t|$ by Lemma 2, and recall that since $\mathcal{D}$ violates (T4$'$) we have $|X| < k \leqslant |V_t|$ for all $t \in t_1 T t_2$. Thus (1) implies that for every $h > |X|$ the number of parts of order h is no greater in $\mathcal{D}'$ than it was in $\mathcal{D}$, while (2) implies that for some such h this number has gone down.

For the proof of (1) let $t \in T$ be given, assume for notational simplicity that $i = 1$, and suppose that $|W_t^1| = |V_t|$; we show that $W_t^2 \subseteq X$. If not then V_t meets C_2, and for every vertex of V_t in C_2 some $x \in X$ was included in W_t^1 when it was formed from V_t. Let Y be the set of all those x:

$$Y := W_t^1 \setminus V_t = \{\, x \in X \mid t \in t_x T t_2 \,\} \setminus V_t, \tag{$*$}$$

and put

$$X' := (X \setminus Y) \cup (V_t \cap C_2)$$

(Fig. 1). As indicated above, $|W_t^1| = |V_t|$ implies that $|Y| = |V_t \cap V(C_2)|$, so $|X'| = |X| < k$. Our aim is to show that X' should have been chosen instead of X for our minimal counterexample.


FIGURE 1.   Obtain X' from X by replacing the vertices of Y with vertices from $V_t \cap V(C)$.

To this end, let us show that X' separates both Z_1 and V_t from Z_2 in G. Any path P in $G - X'$ from either Z_1 or V_t to Z_2 has a last vertex y in Y. Then $\mathring{y}P \subseteq C_2$, and yP contains a V_{t_y}–V_{t_2} path. Since $t \in t_y T t_2$ by ($*$), yP meets V_t (Lemma 1), and as $y \notin V_t$ it must do so in C_2, ie. in X'.

If $t \in \mathring{t_1}Tt_2$ then the above implies that for any k-set $Z \subseteq V_t$ the quadruple (t, t_2, Z, Z_2) violates (T4′), which contradicts the choice of t_1 and t_2. So $t \notin \mathring{t_1}Tt_2$. We complete the proof of (1) by showing that $d_{x'} < d_y$ for all $x' \in X' \smallsetminus X$ and all $y \in Y$ ($= X \smallsetminus X'$); then X' should have been chosen instead of X for our counterexample (t_1, t_2, Z_1, Z_2).

Let x' and y be given. By $(\ast)$ t separates t_y from t_2 in T, and hence as $t \notin \mathring{t_1}Tt_2$ from the whole path t_1Tt_2. Writing d for the distance of t from t_1Tt_2 in T, we thus obtain $d_{x'} \leqslant d < d_y$ by $x' \in V_t$ and the definition of t_y.

For the proof of (2) it suffices by (1) to find $t \in t_1Tt_2$ such that V_t crosses X, ie. meets both C_1 and C_2. Suppose there is no such t. Since V_{t_1} meets C_1 and V_{t_2} meets C_2, this means that t_1Tt_2 has an edge tt' such that $V_t \subseteq H_1$ and $V_{t'} \subseteq H_2$. But then $V_t \cap V_{t'} \subseteq X$ and tt' satisfies (T4′), a contradiction. $\quad\square$

4. The tree-width duality theorem

Let G be a graph. Two subsets of $V(G)$ are said to *touch* if they have a vertex in common or G contains an edge between them. Following [6], we call a set $\mathcal{B}$ of mutually touching connected vertex sets a *bramble*. A subset of $V(G)$ is said to *cover* $\mathcal{B}$ if it meets every element of $\mathcal{B}$. The least number of vertices covering a bramble is its *order*. A typical example of an order n bramble is the set of crosses in the $n \times n$ grid; see [2].

Lemma 4. *Any set of vertices separating two covers of a bramble also covers that bramble.*

Proof. Since each set in the bramble is connected and meets both of the covers, it also meets any set separating these covers. $\quad\square$

Theorem 5. *Let $k \geqslant 0$ be an integer. A graph has tree-width $\geqslant k$ if and only if it contains a bramble of order $> k$.*

Proof. For the backward implication, let $\mathcal{B}$ be any bramble in a graph G. We show that every tree-decomposition $(V_t)_{t \in T}$ of G has a part that covers $\mathcal{B}$. For every edge $e \in T$, at least one of the two components T' of $T - e$ is such that $\bigcup_{t \in T'} V_t$ covers $\mathcal{B}$, because the sets in $\mathcal{B}$ are connected and touch (cf. Lemma 1); we then orient the edge e towards T'. Having oriented every edge of T we let s be the last vertex of a maximal directed path in T and note that V_s covers $\mathcal{B}$.

To prove the forward direction, we now assume that G contains no bramble of order $> k$. We show that for every bramble $\mathcal{B}$ in G there is a $\mathcal{B}$-*admissible* tree-decomposition of G, one in which any part of order $> k$ fails to cover $\mathcal{B}$. For $\mathcal{B} = \emptyset$ this implies that $\text{tw}(G) < k$, because every set covers the empty bramble.

Let $\mathcal{B}$ be given, and assume inductively that for every bramble $\mathcal{B}'$ containing more sets than $\mathcal{B}$ there is a $\mathcal{B}'$-admissible tree-decomposition of G. (The

induction starts, since no bramble in G has more than $2^{|G|}$ sets.) Let $X \subseteq V(G)$ be a cover of $\mathcal{B}$ with as few vertices as possible; then $\ell := |X| \leqslant k$ is the order of $\mathcal{B}$. Our aim is to show the following:

> For every component C of $G - X$ there exists a $\mathcal{B}$-admissible tree-decomposition of $G[C \cup X]$ with X as a part. $\hspace{1em}(*)$

Then these tree-decompositions can be combined to a $\mathcal{B}$-admissible tree-decomposition of G by identifying their nodes corresponding to X. (If $X = V(G)$, then the tree-decomposition with X as its only part is $\mathcal{B}$-admissible.)

So let C be a fixed component of $G - X$, write $H := G[C \cup X]$, and put $\mathcal{B}' := \mathcal{B} \cup \{C\}$. If $\mathcal{B}'$ is not a bramble then C fails to touch some element of $\mathcal{B}$, and hence $Y := V(C) \cup N(C)$ does not cover $\mathcal{B}$. Then the tree-decomposition of H consisting of the two parts X and Y satisfies $(*)$.

So we may assume that $\mathcal{B}'$ is a bramble. Since X covers $\mathcal{B}$, we have $C \notin \mathcal{B}$ and hence $|\mathcal{B}'| > |\mathcal{B}|$. Our induction hypothesis therefore ensures that G has a $\mathcal{B}'$-admissible tree-decomposition $\mathcal{D} = (V_t)_{t \in T}$. If this decomposition is also $\mathcal{B}$-admissible, there is nothing more to show. If not, then one of its parts of order $> k$, V_s say, covers $\mathcal{B}$. Since no set of fewer than ℓ vertices covers $\mathcal{B}$, Lemma 4 implies with Menger's theorem that V_s and X are linked in G by ℓ disjoint paths P_x with $x \in P_x$ for all $x \in X$. As V_s fails to cover $\mathcal{B}'$ and hence lies in $G - C$, so do these paths.

For each $x \in X$ pick a 'home' node $t_x \in T$ with $x \in V_{t_x}$, and consider $\mathcal{D}_s(H) = (W_t)_{t \in T}$ as in Lemma 2. Then $W_s = X$. To complete the proof of $(*)$, we show that $\mathcal{D}_s(H)$ is $\mathcal{B}$-admissible. Consider any W_t of order $> k$. Then W_t meets C, because $W_t \subseteq V(H)$ and $|X| \leqslant k$. Since $\mathcal{D}$ is $\mathcal{B}'$-admissible and $|V_t| \geqslant |W_t| > k$ by Lemma 2, we know that V_t fails to meet some $B \in \mathcal{B}$; let us show that W_t does not meet this B either. If it does, it must do so in some $x \in X$ with $x \in W_t \smallsetminus V_t$. Then B is a connected set meeting both V_s and V_{t_x} but not V_t, contradicting $t \in sTt_x$ by Lemma 1. $\hspace{1em}\square$

Theorem 5 can be restated in terms of the *bramble number* of a graph, the largest order of any bramble in it. It then says that the tree-width of a graph is exactly one less than its bramble number.

References

[1] P. Bellenbaum, *Schlanke Baumzerlegungen von Graphen*, Diplomarbeit Universität Hamburg 2000.

[2] R. Diestel, *Graph theory*, 2nd edition, Springer-Verlag 2000 and
`http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/download.html`

[3] R. Diestel, K.Yu. Gorbunov, T.R. Jensen & C. Thomassen, Highly connected sets and the excluded grid theorem, *J. Combin. Theory B* **75** (1999), 61–73. (Also included in [2] with improved exposition.)

[4] J.F. Geelen, A.M.H. Gerards & G. Whittle, Branch width and well-quasi-ordering in matroids and graphs, *J. Combin. Theory B* **84** (2002), 270–290.

[5] J.F. Geelen, B. Richter, & G. Salazar, Embedding grids in surfaces, preprint 2000.

[6] B.A. Reed, Tree-width and tangles: a new connectivity measure and some applications, in (R.A. Bailey, ed) *Surveys in Combinatorics 1997*, Cambridge University Press 1997, 87–162.

[7] N. Robertson & P.D. Seymour, Graph minors IV; tree-width and well-quasi-ordering, *J. Combin. Theory B* **48** (1990), 227–254.

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

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

[10] C. Thomassen, A simpler proof of the excluded minor theorem for higher surfaces, *J. Combin. Theory B* **70** (1997), 306–311.