Journal of Combinatorial Theory, Series B **75**, 61–73 (1999)
Article ID jctb.1998.1862, available online at http://www.idealibrary.com on IDEAL®

Highly Connected Sets and the Excluded Grid Theorem

Reinhard Diestel and Tommy R. Jensen

Faculty of Mathematics, Technical University of Chemnitz, D-09107 Chemnitz, Germany

Konstantin Yu. Gorbunov

Institute of New Technologies, Nizhnyaya Radishchevskaya 10, Moscow 109004, Russia

and

Carsten Thomassen

Mathematics Institute, Technical University of Denmark, DK-2800 Lyngby, Denmark

Received June 4, 1998

We present a short proof of the excluded grid theorem of Robertson and Seymour, the fact that a graph has no large grid minor if and only if it has small tree-width. We further propose a very simple obstruction to small tree-width inspired by that proof, showing that a graph has small tree-width if and only if it contains no large highly connected set of vertices. © 1999 Academic Press

1. INTRODUCTION

The following theorem of Robertson and Seymour [5] plays a fundamental role in their theory of graph minors:

THEOREM 1. *Given any graph X, the graphs without an X minor have bounded tree-width if and only if X is planar.*

Since planar grids can have arbitrarily large tree-width (see below), the "only if" direction here is immediate: if X is non-planar then no grid has an X minor, and hence the graphs without an X minor have unbounded tree-width. Conversely, we have to show that forbidding any planar minor bounds the tree-width of a graph. And again, since every planar graph G is the minor of some large enough grid (take a drawing of G with "fat"

0095-8956/99 $30.00

vertices and superimpose a drawing of a sufficiently fine grid), it suffices to show the following:

THEOREM 2. *For every integer r there is an integer k such that every graph of tree-width at least k has an $r \times r$ grid minor.*

Proofs of Theorem 2 have been given by Robertson and Seymour [5], by Robertson, Seymour, and Thomas [7], and by Reed [3]. All these proofs are long and technical. Our main purpose in this paper is to offer a short and self-contained new proof of Theorem 2. This will be given in Section 3, which can be read independently of the rest of the paper.

We remark that our proof of Theorem 2 may be combined with [4] and [8] to give the shortest known proof of one of the main corollaries of the Robertson–Seymour graph minor theorem ("Wagner's conjecture"), the "generalized Kuratowski" result that the graphs embeddable in any fixed surface are characterized by finitely many forbidden minors. A proof of the graph minor theorem itself is sketched in [1, Chapter 12]; among other things, the sketch indicates the role that Theorem 1 plays in that proof.

Our second aim in this paper is to draw attention to another obstruction to small tree-width, implied by (but different from) large grid minors: large "highly connected" sets of vertices. In Section 2 we give a very simple proof that a graph has small tree-width if and only if it contains no such set of vertices. A result with a similar flavour has been obtained by Reed as a spin-off of the theory of "brambles" [3, combine Lemma 3.4 with Theorem 2.11].

Our terminology follows [1]. (A general introduction to tree-decompositions and graph minors may also be found there, as well as in [9].) The vertex sets into which a tree-decomposition decomposes a graph will be called the *parts* of that decomposition. For the notion of tree-width, we just recall that tree-decompositions of width $<k$ may have parts containing up to k vertices; thus, trees have tree-width 1. If C is a subgraph of a graph G, we write $N(C)$ for its set of *neighbours* in $G-C$, the set of vertices in $G-C$ adjacent to a vertex in C. A *separation* of G is an ordered pair (A, B) of subgraphs of G such that $A \cup B = G$ and $E(A) \cap E(B) = \emptyset$; its *order* is the number $|A \cap B|$. The *$n \times n$ grid* is the graph on $\{1, ..., n\}^2$ with edge set $\{(i, j)(i', j'): |i-i'| + |j-j'| = 1\}$. We call a set $X \subseteq V(G)$ *k-connected* in G if $|X| \geqslant k$ and for all subsets $Y, Z \subseteq X$ with $|Y| = |Z| \leqslant k$ there are $|Y|$ disjoint $Y-Z$ paths in G. (The sets Y and Z are not required to be disjoint.) X is *externally k-connected* if, in addition, the required paths can be chosen without an inner vertex or edge in $G[X]$. For example, the vertex set of any k-connected subgraph of G is k-connected in G (though not necessarily externally), but also any horizontal path of the $k \times k$ grid is k-connected in the grid, even externally.

2. HIGHLY CONNECTED SETS

In this section we show that a graph has small tree-width if and only if it has no large highly connected sets of vertices. The proof of the first part of the following proposition uses no more than standard tree-decomposition techniques; we include it for the convenience of those readers new to the subject.

PROPOSITION 3. *Let G be a graph and $k > 0$ an integer.*

(i) *If G has tree-width $< k$ then G contains no $(k+1)$-connected set of size $\geqslant 3k$.*

(ii) *Conversely, if G contains no externally $(k+1)$-connected set of size $\geqslant 3k$ then G has tree-width $< 4k$.*

Proof. (i) Choose a tree-decomposition $(T, (V_t)_{t \in T})$ of G of width $< k$, without loss of generality so that none of the parts V_t is contained in another. Then for every edge $e = rs$ of T, the set $V_r \cap V_s$ has $< k$ vertices and separates the sets $U_r := \bigcup_{t \in T_r} V_t$ and $U_s := \bigcup_{t \in T_s} V_t$ in G; here, T_r and T_s denote the components of $T - e$ containing r and s, respectively. Any separation (A, B) of G with vertex sets $\{V(A), V(B)\} = \{U_r, U_s\}$ (and hence $V(A \cap B) = V_r \cap V_s$) is said to *correspond to e*.

Suppose X is a $(k+1)$-connected set of size $\geqslant 3k$ in G. Orient every edge e of T towards the component T' of $T - e$ for which $|X \cap \cup_{t \in T'} V_t|$ is greater, breaking ties arbitrarily. Choose a vertex $t \in T$ so that all the edges $e_1, ..., e_n$ of T at t point towards t. For every $i = 1, ..., n$ pick a separation (A_i, B_i) corresponding to e_i, with $V_t \subseteq B_i$. Then $|V(A_i) \cap X| < k$: otherwise, both A_i and B_i would have k vertices in X, and we could extend $V(A_i \cap B_i) \cap X$ to k-subsets $Y \subseteq V(A_i) \cap X$ and $Z \subseteq V(B_i) \cap X$ that cannot be linked by k disjoint paths in G (since $|A_i \cap B_i| < k$).

Now let $i \leqslant n$ be minimal such that $|V(A_1 \cup \cdots \cup A_i) \cap X| > k$, and put $A := A_1 \cup \cdots \cup A_i$ and $B := B_1 \cap \cdots \cap B_i$. By the minimality of i and since $|V(A_i) \cap X| < k$, we have $|V(A) \cap X| < 2k$, so $|V(B) \cap X| > |X| - 2k \geqslant k$. As before, we may extend $V(A \cap B) \cap X$ to $(k+1)$-sets $Y \subseteq V(A) \cap X$ and $Z \subseteq V(B) \cap X$. As V_t separates these sets in G and $|V_t| \leqslant k$, this contradicts our assumption that X is $(k+1)$-connected in G.

(ii) We prove the following more general assertion:

> *If $h \geqslant k$ and G contains no externally k-connected set of size h, then G has tree-width $< h + k - 1$.*

Let $U \subseteq V(G)$ be maximal such that $G[U]$ has a tree-decomposition $\mathscr{D}$ of width $< h + k - 1$ such that every component C of $G - U$ has at most h neighbours in U and these lie in one part of $\mathscr{D}$ (depending on C).

We claim that $U = V(G)$. Suppose not. Let C be a component of $G - U$ and write $X := N(C)$. By assumption, $|X| \leqslant h$. In fact $|X| = h$, since otherwise for any $v \in V(C)$ we could add $X \cup \{v\}$ to $\mathscr{D}$ as a new part, contradicting the maximality of U. Hence by assumption, X is not externally k-connected in G; let $Y, Z \subseteq X$ be sets to witness this.

By Menger's theorem, Y and Z are separated in $H := G[V(C) \cup Y \cup Z] - E(G[Y \cup Z])$ by a set S of fewer than $|Y| = |Z| \leqslant k$ vertices. Let $X_Y := (X \backslash Z) \cup S$ and $X_Z := (X \backslash Y) \cup S$. Clearly, $|X \cup S| \leqslant h + k - 1$ and $|X_Y|, |X_Z| < |X| = h$. Moreover, any component $C' \subseteq C$ of $G - (U \cup S)$ has all its neighbours in $X \cup S$, and hence either in X_Y or in X_Z: otherwise $H - S$ would contain a $Y - Z$ path through C'.

Extending U to $U \cup S$ and adding $X \cup S$ to $\mathscr{D}$ as a new part, we obtain a contradiction to the maximality of U. (Note that $S \cap C \neq \emptyset$, since $|S| < |Y| = |Z|$ and $Y, Z \subseteq N(C)$.) ∎

It is perhaps interesting to note that Proposition 3(i) is best possible (or nearly so) in various ways, and remains so even if we weaken its assertion by inserting "externally" before "$(k+1)$-connected." For example, a complete bipartite graph with vertex sets X and Y of sizes $k-1$ and $n \geqslant k-1$, respectively, has tree-width $<k$ (with parts $X \cup \{y\}$, $y \in Y$) and clearly Y is an externally $(k-1)$-connected set whose size is not bounded as a function of k. Only slightly less trivially, consider the complete bipartite graph G with vertex sets X and Y of sizes $k \geqslant 5$ and $3k$, respectively. Partition Y into three sets Y_1, Y_2, Y_3 of size k, and delete a perfect matching from each of the three bipartite subgraphs $G[X \cup Y_i]$. The resulting graph again has tree-width $<k$ (with parts X and $(X \backslash \{x\}) \cup \{y\}$ for all $y \in Y$, where x is the unique non-neighbour of y), and an easy application of Hall's theorem shows that Y is externally k-connected in it.

The above example shows that the value of k in the premise and the value of $(k+1)$ in the conclusion of Proposition 3(i) are best possible. The value of $3k$ in the conclusion is also essentially best possible. This is exemplified by the following graph G of tree-width $<k$ that contains an externally $(k+1)$-connected (even $|X|$-connected) set X of size $3(k-1)$. First we define a partially ordered set, as follows. Starting with a linearly ordered $(k-1)$-set R, we put three linearly ordered $(k-1)$-sets S_1, S_2, S_3 above R, letting elements from different S_i be incomparable. To make this into a graph, we add all possible edges on R and join every vertex from one of the S_i to all the vertices in the $(k-1)$-chain directly below it. (Note that this graph has a tree-decomposition into its k-cliques.) Finally, we add new independent $(k-1)$-sets X_1, X_2, X_3, joining X_i completely to S_i for each $i = 1, 2, 3$. It is now easily checked that any two (disjoint) sets $Y, Z \subseteq X_1 \cup X_2 \cup X_3$ of equal size can be linked in G by $|Y| = |Z|$ disjoint paths: vertices y, z in the same X_i can be linked via the corresponding S_i using

its highest vertices, and any remaining pairs y, z can be joined via the lower vertices of the S_i and R.

Proposition 3(ii) is also best possible: the complete graph on $4k$ vertices shows that we cannot strengthen the conclusion, while the complete graph on $4k+1$ vertices shows that we cannot weaken the premise in either way. Similarly, K_{h+k-1} and K_{h+k} (for $h \geqslant 2k+2$) show that the more general assertion we prove is best possible.

3. GRID MINORS

We now present our proof of Theorem 2. Very roughly, we shall assume that a given graph G has large tree-width, find a large highly connected set X in G as in Proposition 3(ii), and use its connecting paths P to form a grid. Of course, this will be possible only if those paths intersect sufficiently. If they do not, we shall try instead to partition X into many sets that can be linked pairwise by mutually disjoint paths, so that contracting these sets will give us a subdivision of a large complete graph. Since we may only contract connected sets when forming a minor, our first task will thus be to strengthen Proposition 3(ii) so as to give X a partition into many sets that can be made connected in a part of G not used by the paths P.

In order to make this section self-contained, we prove all the lemmas that we need from first principles. This goes in particular for our first lemma, the strengthening of Proposition 3(ii) indicated above.

Let us call a separation (A, B) a *premesh* if all the edges of $A \cap B$ lie in A and A contains a tree T with the following properties:

- T has maximum degree $\leqslant 3$;
- every vertex of $A \cap B$ lies in T and has degree $\leqslant 2$ in T;
- T has a leaf in $A \cap B$ (that is, a vertex of degree $\leqslant 1$).

A premesh (A, B) will be called a *k-mesh* if $V(A \cap B)$ is externally k-connected in B, and the graph $G = A \cup B$ is said to *have* this premesh or k-mesh.

LEMMA 4. *Let G be a graph and let $h \geqslant k \geqslant 1$ be integers. If G has no k-mesh of order h then G has tree-width $< h+k-1$.*

Proof. We may assume that G is connected. Let $U \subseteq V(G)$ be maximal such that $G[U]$ has a tree-decomposition $\mathscr{D}$ of width $< h+k-1$, with the additional property that, for every component C of $G-U$, the neighbours of C in U lie in one part of $\mathscr{D}$ and $(G-C, \bar{C})$ is a premesh of order $\leqslant h$, where $\bar{C} := G[V(C) \cup N(C)] - E(G[N(C)])$. Clearly, $U \neq \emptyset$.

We claim that $U = V(G)$. Suppose not. Let C be a component of $G - U$, put $X := N(C)$, and let T be a tree associated with the premesh $(G - C, \bar{C})$.

By assumption, $|X| \leqslant h$; let us show that equality holds here. If not, let $u \in X$ be a leaf of T and v a neighbour of u in C. Put $U' := U \cup \{v\}$ and $X' := X \cup \{v\}$, let T' be the tree obtained from T by joining v to u, and let $\mathscr{D}'$ be the tree-decomposition of $G[U']$ obtained from $\mathscr{D}$ by adding X' as a new part. Clearly $\mathscr{D}'$ still has width $< h + k - 1$. Consider a component C' of $G - U'$. If $C' \cap C = \emptyset$ then C' is also a component of $G - U$, so $N(C')$ lies inside a part of $\mathscr{D}$ (and hence of $\mathscr{D}'$), and $(G - C', \bar{C}')$ is a premesh of order $\leqslant h$ by assumption. If $C' \cap C \neq \emptyset$, then $C' \subseteq C$ and $N(C') \subseteq X'$. Moreover, $v \in N(C')$: otherwise $N(C') \subseteq X$ would separate C' from v, contradicting the fact that C' and v lie in the same component C of $G - X$. Since v is a leaf of T', it is straightforward to check that $(G - C', \bar{C}')$ is again a premesh of order $\leqslant h$, contrary to the maximality of U.

Thus $|X| = h$, so by assumption our premesh $(G - C, \bar{C})$ cannot be a k-mesh. Hence by Menger's theorem, there are sets $Y, Z \subseteq X$ of equal size $\leqslant k$ that are separated in $H := \bar{C}[V(C) \cup Y \cup Z]$ by a set S of $k' < |Y| = |Z|$ vertices, one from each of a family $(P_s \mid s \in S)$ of disjoint $Y - Z$ paths in H. Put $X' := X \cup S$ and $U' := U \cup S$, and let $\mathscr{D}'$ be the tree-decomposition of $G[U']$ obtained from $\mathscr{D}$ by adding X' as a new part. Clearly, $|X'| \leqslant |X| + |S| \leqslant h + k - 1$. We show that U' contradicts the maximality of U.

Since $Y \cup Z \subseteq N(C)$ and $|S| < |Y| = |Z|$ we have $S \cap C \neq \emptyset$, so U' is larger than U. Let C' be a component of $G - U'$. If $C' \cap C = \emptyset$, we argue as earlier. So $C' \subseteq C$ and $N(C') \subseteq X'$. As before, at least one neighbour v of C' lies in $S \cap C$. By definition of S, C' cannot have neighbours in both $Y \backslash S$ and $Z \backslash S$; we assume it has none in $Y \backslash S$. Let T' be the union of T and all the $Y - S$ subpaths of paths P_s with $s \in N(C') \cap C$; since these subpaths start in $Y \backslash S$ and have no inner vertices in X', they cannot meet C'. Therefore $(G - C', \bar{C}')$ is a premesh with tree T' and leaf v; the degree conditions on T' are easily checked. Its order is $|N(C')| \leqslant |X| - |Y| + |S| = h - |Y| + k' < h$, a contradiction to the maximality of U. ∎

Lemma 5. *Let $k \geqslant 2$ be an integer. Let T be a tree of maximum degree $\leqslant 3$ and $X \subseteq V(T)$. Then T has a set E of edges such that every component of $T - E$ has between k and $2k - 2$ vertices in X, except that one such component may have fewer vertices in X.*

Proof. Induction on $|X|$. If $|X| \leqslant 2k - 2$ we put $E = \emptyset$. So assume that $|X| \geqslant 2k - 1$. Let e be an edge of T such that some component T' of $T - e$ has at least k vertices in X and T' is as small as possible. As $\Delta(T) \leqslant 3$, the end of e in T' has degree at most two in T', so the minimality of T' implies that $|X \cap V(T')| \leqslant 2k - 2$. We finish by applying the induction hypothesis to $T - T'$. ∎

Lemma 6. *Let G be a bipartite graph with bipartition (A, B), $|A| = a$, $|B| = b$, and let $c \leqslant a$ and $d \leqslant b$ be positive integers. Assume that G has at most $(a-c)(b-d)/d$ edges. Then there exist $C \subseteq A$ and $D \subseteq B$ such that $|C| = c$ and $|D| = d$ and $C \cup D$ is independent in G.*

Proof. As $\|G\| \leqslant (a-c)(b-d)/d$, fewer than $b-d$ vertices in B have more than $(a-c)/d$ neighbours in A. Choose $D \subseteq B$ so that $|D| = d$ and each vertex in D has at most $(a-c)/d$ neighbours in A. Then D sends a total of at most $a-c$ edges to A, so A has a subset C of c vertices without a neighbour in D. ∎

Given a tree T, call an r-tuple $(x_1, ..., x_r)$ of distinct vertices of T *good* if, for every $j = 1, ..., r-1$, the $x_j - x_{j+1}$ path in T contains none of the other vertices in this r-tuple.

Lemma 7. *Every tree of order $\geqslant r(r-1)$ has a good r-tuple of vertices.*

Proof. If x is any vertex of a tree T, then T is the union of its subpaths $x \cdots y$, where y ranges over its leaves. Hence unless one of these paths has at least r vertices, T has at least $|T|/(r-1)$ leaves. Since any path of r vertices and any set of r leaves defines a good r-tuple in T, this proves the assertion. ∎

Our next lemma shows how to obtain a grid from two large systems of paths that intersect in a particularly orderly way.

Lemma 8. *Let $d, r \geqslant 2$ be integers such that $d \geqslant r^{2r+2}$. Let G be a graph containing a set $\mathscr{H}$ of $r^2 - 1$ disjoint paths and a set $\mathscr{V} = \{V_1, ..., V_d\}$ of d disjoint paths. Assume that every path in $\mathscr{V}$ meets every path in $\mathscr{H}$, and that each path $H \in \mathscr{H}$ consists of d consecutive (vertex disjoint) segments such that V_i meets H only in its ith segment, for every $i = 1, ..., d$. Then G has an $r \times r$ grid minor.*

Proof. For each $i = 1, ..., d$, consider the graph with vertex set $\mathscr{H}$ in which two paths are adjacent whenever V_i contains a subpath between them that meets no other path in $\mathscr{H}$. Since V_i meets every path in $\mathscr{H}$, this is a connected graph; let T_i be a spanning tree in it. Since $|\mathscr{H}| \geqslant r(r-1)$, Lemma 7 implies that each of these $d \geqslant r^2(r^2)^r$ trees T_i has a good r-tuple of vertices. Since there are no more than $(r^2)^r$ distinct r-tuples on $\mathscr{H}$, some r^2 of the trees T_i have a common good r-tuple $(H^1, ..., H^r)$. Let $I = \{i_1, ..., i_{r^2}\}$ be the index set of these trees (with $i_j < i_k$ for $j < k$) and put $\mathscr{H}' := \{H^1, ..., H^r\}$.

Here is an informal description of how we construct our $r \times r$ grid. Its "horizontal" paths will be the paths $H^1, ..., H^r$. Its "vertical" paths will be pieced together edge by edge, as follows. The $r-1$ edges of the first vertical path will come from the first $r-1$ trees T_i, trees with their index i among the first r elements of I. More precisely, its "edge" between H^j and H^{j+1} will be the sequence of subpaths of V_{i_j} (together with some connecting horizontal bits taken from paths in $\mathscr{H}\setminus\mathscr{H}'$) induced by the edges of an H^j-H^{j+1} path in T_{i_j} that has no inner vertices in $\mathscr{H}'$. (This is why we need $(H^1, ..., H^r)$ to be a good r-tuple in every tree T_i.) Similarly, the jth edge of the second vertical path will come from an H^j-H^{j+1} path in $T_{i_{r+j}}$, and so on. To merge these individual edges into r vertical paths, we then contract in each H^j the initial segment that meets the first r paths V_i with $i \in I$, then contract the segment that meets the following r paths V_i with $i \in I$, and so on.

Formally, we proceed as follows. For all $j, k \in \{1, ..., r\}$, consider the minimal subpath H^j_k of H^j that contains the ith segment of H^j for all i with $i_{(k-1)r} < i \leqslant i_{kr}$ (put $i_0 := 0$). Let $\hat{H}^j$ be obtained from H^j by first deleting any vertices following its i_{r^2}th segment and then contracting every subpath H^j_k to one vertex v^j_k. Thus, $\hat{H}^j = v^j_1 ... v^j_r$.

Given $j \in \{1, ..., r-1\}$ and $k \in \{1, ..., r\}$, we have to define a path V^j_k that will form the subdivided "vertical edge" $v^j_k v^{j+1}_k$. This path will consist of segments of the path V_i together with some otherwise unused segments of paths from $\mathscr{H}\setminus\mathscr{H}'$, for $i := i_{(k-1)r+j}$; recall that, by definition of $\hat{H}^j$ and $\hat{H}^{j+1}$, this V_i does indeed meet H^j and H^{j+1} precisely in vertices that were contracted into v^j_k and v^{j+1}_k, respectively. To define V^j_k, consider an H^j-H^{j+1} path $P = H_1 \cdots H_t$ in T_i that has no inner vertices in $\mathscr{H}'$. Every edge $H_s H_{s+1}$ of P corresponds to an $H_s - H_{s+1}$ subpath of V_i that has no inner vertex on any path in $\mathscr{H}$. Together with (parts of) the ith segments of $H_2, ..., H_{t-1}$, these subpaths of V_i form an H^j-H^{j+1} path P' that has no inner vertices on any of the paths $H^1, ..., H^r$ and meets no path from $\mathscr{H}$ outside its ith segment. Replacing the ends of P' on H^j and H^{j+1} with v^j_k and v^{j+1}_k, respectively, we obtain our desired path V^j_k forming the jth (subdivided) edge of the kth "vertical" path of our grid. Since the paths P' are disjoint for different i and different pairs (j, k) do give rise to different i, the paths V^j_k are disjoint except for possible common ends v^j_k. Moreover, they have no inner vertices on any of the paths $H^1, ..., H^r$, because none of these H^j is an inner vertex of any of the paths $P \subseteq T_i$ used in the construction of V^j_k. ∎

We are now ready to prove the following quantitative version of Theorem 2.

Theorem 9. *Let $r, m > 0$ be integers, and let G be a graph of tree-width at least $r^{4m^2(r+2)}$. Then G contains either K_m or the $r \times r$ grid as a minor.*

Proof. Since K_{r^2} contains the $r \times r$ grid as a subgraph, we may assume that $2 \leqslant m \leqslant r^2$. Put $c := r^{4(r+2)}$ and $k := c^{2\binom{m}{2}}$. Then $2m+2 \leqslant c$, so G has tree-width at least $c^{m^2} \geqslant (2m+2)\,k$: more than enough for Lemma 4 to ensure that G contains a k-mesh (A, B) of order $(2m+1)(k-1)$. Let $T \subseteq A$ be a tree associated with the premesh (A, B); thus, $X := V(A \cap B) \subseteq V(T)$. By Lemma 5, T has $(|X| - (k-1))/(2k-2) = m$ disjoint subtrees each containing at least k vertices of X; let $A_1, ..., A_m$ be the vertex sets of these trees. By definition of a k-mesh, B contains for all $1 \leqslant i < j \leqslant m$ a set $\mathscr{P}_{ij}$ of k disjoint $A_i - A_j$ paths that have no inner vertices in A. These sets $\mathscr{P}_{ij}$ will shrink a little and be otherwise modified later in the proof, but they will always consist of "many" disjoint $A_i - A_j$ paths.

One option in our proof will be to find single paths $P_{ij} \in \mathscr{P}_{ij}$ that are disjoint for different pairs ij and thus link up the sets A_i to form a K_m minor of G. If this fails, we shall instead exhibit two specific sets $\mathscr{P}_{ij}$ and $\mathscr{P}_{pq}$ such that many paths of $\mathscr{P}_{ij}$ meet many paths of $\mathscr{P}_{pq}$, forming an $r \times r$ grid between them by Lemma 8.

Let us impose a linear ordering on the index pairs ij by fixing an arbitrary bijection σ: $\{ij \mid 1 \leqslant i < j \leqslant m\} \to \{0, 1, ..., \binom{m}{2} - 1\}$. For $\ell = 0, 1, ...$ in turn, we shall consider the pair pq with $\sigma(pq) = \ell$ and choose an $A_p - A_q$ path P_{pq} that is disjoint from all previously selected such paths, i.e., from the paths P_{st} with $\sigma(st) < \ell$. At the same time, we shall replace all the "later" sets $\mathscr{P}_{ij}$—or what has become of them—by smaller sets containing only paths that are disjoint from P_{pq}. Thus for each pair ij, we shall define a sequence $\mathscr{P}_{ij} = \mathscr{P}^0_{ij}, \mathscr{P}^1_{ij}, ...$ of smaller and smaller sets of paths, which eventually collapses to $\mathscr{P}^\ell_{ij} = \{P_{ij}\}$ when ℓ has risen to $\ell = \sigma(ij)$.

More formally, let $\ell^* \leqslant \binom{m}{2}$ be maximal such that, for all $0 \leqslant \ell < \ell^*$ and all $1 \leqslant i < j \leqslant m$, there exist sets $\mathscr{P}^\ell_{ij}$ satisfying the following five conditions:

(i) $\mathscr{P}^\ell_{ij}$ is a non-empty set of disjoint $A_i - A_j$ paths in B that meet A only in their endpoints.

As soon as a set $\mathscr{P}^\ell_{ij}$ is defined, we shall write $H^\ell_{ij} := \bigcup \mathscr{P}^\ell_{ij}$ for the union of its paths.

(ii) If $\sigma(ij) < \ell$ then $\mathscr{P}^\ell_{ij}$ has exactly one element P_{ij}, and P_{ij} does not meet any path belonging to a set $\mathscr{P}^\ell_{st}$ with $ij \neq st$.

(iii) If $\sigma(ij) = \ell$, then $|\mathscr{P}^\ell_{ij}| = k/c^{2\ell}$.

(iv) If $\sigma(ij) > \ell$, then $|\mathscr{P}^\ell_{ij}| = k/c^{2\ell+1}$.

(v) If $\ell = \sigma(pq) < \sigma(ij)$, then for every $e \in E(H^\ell_{ij}) \backslash E(H^\ell_{pq})$ there are no $k/c^{2\ell+1}$ disjoint paths from A_i to A_j in the graph $(H^\ell_{pq} \cup H^\ell_{ij}) - e$.

Note that, since $\sigma(ij) < \binom{m}{2}$ by definition of σ, conditions (iii) and (iv) imply that $|\mathscr{P}^\ell_{ij}| \geqslant c^2$ whenever $\sigma(ij) \geqslant \ell$.

Clearly if $\ell^* = \binom{m}{2}$, then by (i) and (ii) we have a (subdivided) K_m minor with branch sets $A_1, ..., A_m$ in G. Suppose then that $\ell^* < \binom{m}{2}$. Let us show that $\ell^* > 0$. Let $pq := \sigma^{-1}(0)$ and put $\mathscr{P}^0_{pq} := \mathscr{P}_{pq}$. To define $\mathscr{P}^0_{ij}$ for $\sigma(ij) > 0$

put $H_{ij} := \bigcup \mathcal{P}_{ij}$, and let $F \subseteq E(H_{ij}) \setminus E(H^0_{pq})$ be maximal such that $(H^0_{pq} \cup H_{ij}) - F$ still contains k/c disjoint paths from A_i to A_j; then let $\mathcal{P}^0_{ij}$ be such a set of paths. As any vertex of A on these paths lies in $A_i \cup A_j$ (by definition of H^0_{pq} and H_{ij}), we may assume that they have no inner vertices in A. Thus our choice of $\mathcal{P}^0_{ij}$ satisfies (i)–(v).

Having shown that $\ell^* > 0$, let us now consider $\ell := \ell^* - 1$. Thus, conditions (i)–(v) are satisfied for ℓ but cannot be satisfied for $\ell + 1$. Let $pq := \sigma^{-1}(\ell)$. If $\mathcal{P}^\ell_{pq}$ contains a path P that avoids a set $\mathcal{Q}_{ij}$ of some $|\mathcal{P}^\ell_{ij}|/c$ of the paths in $\mathcal{P}^\ell_{ij}$ for all ij with $\sigma(ij) > \ell$, then we can define $\mathcal{P}^{\ell+1}_{ij}$ for all ij as before (with a contradiction). Indeed, let $st := \sigma^{-1}(\ell + 1)$ and put $\mathcal{P}^{\ell+1}_{st} := \mathcal{Q}_{st}$. For $\sigma(ij) > \ell + 1$ write $H_{ij} := \bigcup \mathcal{Q}_{ij}$, let $F \subseteq E(H_{ij}) \setminus E(H^{\ell+1}_{st})$ be maximal such that $(H^{\ell+1}_{st} \cup H_{ij}) - F$ still contains at least $|\mathcal{P}^\ell_{ij}|/c^2$ disjoint paths from A_i to A_j, and let $\mathcal{P}^{\ell+1}_{ij}$ be such a set of paths. Setting $\mathcal{P}^{\ell+1}_{pq} := \{P\}$ and $\mathcal{P}^{\ell+1}_{ij} := \mathcal{P}^\ell_{ij} = \{P^\ell_{ij}\}$ for $\sigma(ij) < \ell$ then gives us a family of sets $\mathcal{P}^{\ell+1}_{ij}$ that contradicts the maximality of ℓ^*.

Thus for every path $P \in \mathcal{P}^\ell_{pq}$ there exists a pair ij with $\sigma(ij) > \ell$ such that P avoids fewer than $|\mathcal{P}^\ell_{ij}|/c$ of the paths in $\mathcal{P}^\ell_{ij}$. For some $\lceil |\mathcal{P}^\ell_{pq}|/\binom{m}{2} \rceil$ of these P that pair ij will be the same; let $\mathcal{P}$ denote the set of those P, and keep ij fixed from now on. Note that $|\mathcal{P}| \geqslant |\mathcal{P}^\ell_{pq}|/\binom{m}{2} = c\,|\mathcal{P}^\ell_{ij}|/\binom{m}{2}$ by (iii) and (iv).

Let us use Lemma 6 to find sets $\mathcal{V} \subseteq \mathcal{P} \subseteq \mathcal{P}^\ell_{pq}$ and $\mathcal{H} \subseteq \mathcal{P}^\ell_{ij}$ such that

$$|\mathcal{V}| \geqslant \tfrac{1}{2}|\mathcal{P}| \qquad \left(\geqslant \frac{c}{m^2}\,|\mathcal{P}^\ell_{ij}| \right)$$

$$|\mathcal{H}| = r^2$$

and every path in $\mathcal{V}$ meets every path in $\mathcal{H}$. We have to check that the bipartite graph with vertex sets $\mathcal{P}$ and $\mathcal{P}^\ell_{ij}$ in which $P \in \mathcal{P}$ is adjacent to $Q \in \mathcal{P}^\ell_{ij}$ whenever $P \cap Q = \emptyset$ does not have too many edges. Since every $P \in \mathcal{P}$ has fewer than $|\mathcal{P}^\ell_{ij}|/c$ neighbours (by definition of $\mathcal{P}$), this graph has indeed at most

$$\begin{aligned} |\mathcal{P}|\,|\mathcal{P}^\ell_{ij}|/c &\leqslant |\mathcal{P}|\,|\mathcal{P}^\ell_{ij}|/6r^2 \\ &\leqslant \lfloor |\mathcal{P}|/2 \rfloor\,|\mathcal{P}^\ell_{ij}|/2r^2 \\ &\leqslant \lfloor |\mathcal{P}|/2 \rfloor\,(|\mathcal{P}^\ell_{ij}|/r^2 - 1) \\ &= (|\mathcal{P}| - \lceil |\mathcal{P}|/2 \rceil)(|\mathcal{P}^\ell_{ij}| - r^2)/r^2 \end{aligned}$$

edges, as required. Hence, $\mathcal{V}$ and $\mathcal{H}$ exist as claimed.

Pick a path $Q \in \mathcal{H}$, and put

$$d := \lfloor \sqrt{c}/m \rfloor = \lfloor r^{2r+4}/m \rfloor \geqslant r^{2r+2}.$$

For $n = 1, 2, ..., d-1$ let e_n be the first edge of Q (on its way from A_i to A_j) such that the initial component Q_n of $Q - e_n$ meets at least $nd\,|\mathscr{P}^{\ell}_{ij}|$ different paths from $\mathscr{V}$, and such that e_n is not an edge of H^{ℓ}_{pq}. As any two vertices of Q that lie on different paths from $\mathscr{V}$ are separated in Q by an edge not in H^{ℓ}_{pq}, each of these Q_n meets exactly $nd\,|\mathscr{P}^{\ell}_{ij}|$ paths from $\mathscr{V}$. Put $Q_0 := \varnothing$ and $Q_d := Q$. Since $|\mathscr{V}| \geqslant d^2\,|\mathscr{P}^{\ell}_{ij}|$, we have thus divided Q into d consecutive disjoint segments $Q'_n := Q_n - Q_{n-1}$ $(n = 1, ..., d)$ each meeting at least $d\,|\mathscr{P}^{\ell}_{ij}|$ paths from $\mathscr{V}$.

For each $n = 1, ..., d-1$, Menger's theorem and conditions (iv) and (v) imply that $H^{\ell}_{pq} \cup H^{\ell}_{ij}$ has a set S_n of $|\mathscr{P}^{\ell}_{ij}| - 1$ vertices such that $(H^{\ell}_{pq} \cup H^{\ell}_{ij}) - e_n - S_n$ contains no path from A_i to A_j. Let S denote the union of all these sets S_n. Then $|S| < d\,|\mathscr{P}^{\ell}_{ij}|$, so each Q'_n meets at least one path $V_n \in \mathscr{V}$ that avoids S.

Clearly, each S_n consists of a choice of exactly one vertex x from every path $P \in \mathscr{P}^{\ell}_{ij} \backslash \{Q\}$. Denote the initial component of $P - x$ by P_n, put $P_0 := \varnothing$ and $P_d := P$, and let $P'_n := P_n - P_{n-1}$ for $n = 1, ..., d$. The separation properties of the sets S_n now imply that $V_n \cap P \subseteq P'_n$ for $n = 1, ..., d$ (and hence in particular, that $P'_n \neq \varnothing$, i.e., that $P_{n-1} \subset P_n$). Indeed V_n cannot meet P_{n-1}, because $P_{n-1} \cup V_n \cup (Q - Q_{n-1})$ would then contain an $A_i - A_j$ path in $(H^{\ell}_{pq} \cup H^{\ell}_{ij}) - e_{n-1} - S_{n-1}$, and likewise (consider S_n) V_n cannot meet $P - P_n$. Thus for all $n = 1, ..., d$, the path V_n meets every path $P \in \mathscr{H} \backslash \{Q\}$ precisely in its nth segment P'_n. Applying Lemma 8 to the path systems $\mathscr{H} \backslash \{Q\}$ and $\{V_1, ..., V_d\}$ now yields the desired grid minor. ∎

To conclude, let us remark that our upper bound of $r^{4m^2(r+2)} \leqslant 2^{5r^5 \log r}$ for the tree-width of a graph without an $r \times r$ grid minor is most likely far from best possible. Robertson, Seymour and Thomas [7] obtain an only slightly better bound of about 2^{9r^5}, but they suspect that the correct order might be as low as $r^2 \log r$.

4. HIGHLY CONNECTED SETS AND TANGLES

The proofs of Theorem 2 given in [7] and [3] rely heavily on a concept central to the Robertson–Seymour theory of minors but not so far considered in this paper, the concept of a *tangle* introduced in [6]. Both these proofs build on the fact that, if the tree-width of a graph is large, then the graph contains a large tangle, and use the tangle for the construction of a grid minor.

While it is easy to see that graphs with a large tangle must have large tree-width—the proof is similar to our proof of Proposition 3(i), see [6] or [3]—its converse, the direction needed for the proof of Theorem 2 in [7] and [3], is not immediate. However, this direction follows easily from our Proposition 3(ii), and the purpose of this section is to show how.

Let $k \geqslant 1$ be an integer. A *tangle of order* k in a graph G is a set $\mathscr{T}$ of separations of G, each of order $<k$, that satisfies the following conditions (cf. [6]):

(T1) if (A, B) is any separation of G of order $<k$, then either $(A, B) \in \mathscr{T}$ or $(B, A) \in \mathscr{T}$;

(T2) if $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \mathscr{T}$, then $A_1 \cup A_2 \cup A_3 \neq G$;

(T3) if $(A, B) \in \mathscr{T}$, then $V(A) \neq V(G)$.

As shown in [6, (5.2)], any graph with a tangle of order k has tree-width at least $k-1$, and any graph of tree-width at least $k-1$ has a tangle of order at least $\frac{2}{3}k$. Reed [3] gives a simple proof that graphs of tree-width $\geqslant 3(k-1)$ have a tangle of order k; his proof, however, builds on a non-trivial duality theorem for tree-width due to Seymour and Thomas. Trading just a little more quantitative exactness for simplicity of proof, we observe the following corollary to Proposition 3:

Proposition 10. *Any graph of tree-width at least* $4k$ *has a tangle of order* k.

Proof. If a graph G has tree-width $\geqslant 4k$ then, by Proposition 3(ii), G contains a k-connected set X of size $3k$. Let $\mathscr{T}$ be the set of all separations (A, B) of order $<k$ in G such that $|V(A) \cap X| \leqslant |V(B) \cap X|$. Then $|V(A) \cap X| < k$ for all $(A, B) \in \mathscr{T}$ (since X is k-connected in G but $|A \cap B| < k$), implying (T2) and (T3). Hence $\mathscr{T}$ is a tangle of order k. ∎

REFERENCES

1. R. Diestel, “Graph Theory,” Graduate Texts in Mathematics, Vol. 173, Springer-Verlag, New York, 1997.
2. K. Yu. Gorbunov, “Structural Properties of Context-Free Grammars,” Ph.D. thesis, Moscow State University, 1998; for a partial English translation, see An estimate of the tree-width of a graph which has not a given planar grid as a minor, *in* “1998 Workshop on Graph Theoretic Concepts in Computer Science,” Springer Lecture Notes in Computer Science, Springer-Verlag, Berlin/New York, to appear.
3. B. A. Reed, Tree-width and tangles: A new connectivity measure and some applications, *in* “Surveys in Combinatorics” (R. A. Bailey, Ed.), pp. 87–162, Cambridge Univ. Press, 1997.
4. N. Robertson and P. D. Seymour, Graph Minors. IV. Tree-width and well-quasi-ordering, *J. Combin. Theory Ser. B* **48** (1990), 227–254.
5. N. Robertson and P. D. Seymour, Graph Minors. V. Excluding a planar graph, *J. Combin. Theory Ser. B* **41** (1986), 92–114.
6. N. Robertson and P. D. Seymour, Graph Minors. X. Obstructions to tree-decomposition, *J. Combin. Theory Ser. B* **52** (1991), 153–190.

7. N. Robertson, P. D. Seymour, and R. Thomas, Quickly excluding a planar graph, *J. Combin. Theory Ser. B* **62** (1994), 323–348.
8. C. Thomassen, A simpler proof of the excluded minor theorem for higher surfaces, *J. Combin. Theory Ser. B* **70** (1997), 306–311.
9. C. Thomassen, Embeddings and minors, *in* “Handbook of Combinatorics” (R. L. Graham, M. Grötschel, and L. Lovász, Eds.), Elsevier, Amsterdam, 1996.