A Separator Theorem for Graphs with an Excluded Minor and its Applications

Noga Alon
IBM Almaden Research Center, San Jose, CA 95120 ,USA and
Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel
Paul Seymour
BellCore, 445 South St., Morristown, NJ 07960, USA
Robin Thomas
School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332, USA

Abstract

Let G be an n-vertex graph with nonnegative weights whose sum is 1 assigned to its vertices, and with no minor isomorphic to a given h-vertex graph H. We prove that there is a set X of no more than $h^{3/2}n^{1/2}$ vertices of G whose deletion creates a graph in which the total weight of every connected component is at most $1/2$. This extends significantly a well-known theorem of Lipton and Tarjan for planar graphs. We exhibit an algorithm which finds, given an n-vertex graph G with weights as above and an h-vertex graph H, either such a set X or a minor of G isomorphic to H. The algorithm runs in time $O(h^{1/2}n^{1/2}m)$, where m is the number of edges of G plus the number of its vertices. Our results supply extensions of the many known applications of the Lipton-Tarjan separator theorem from the class of planar graphs (or that of graphs with bounded genus) to any class of graphs with an excluded minor. For example, it follows that for any fixed graph H , given a graph G with n vertices and with no H-minor one can approximate the size of the maximum independent set of G up to a relative error of $1/\sqrt{\log n}$ in polynomial time, find that size exactly and find the chromatic number of G in time $2^{O(\sqrt{n})}$ and solve any sparse system of n linear equations in n unknowns whose sparsity structure

corresponds to G in time $O(n^{3/2})$. We also describe a combinatorial application of our result which relates the tree-width of a graph to the maximum size of a K_h-minor in it.

1   Introduction

A *separation* of a graph G is a pair (A, B) of subsets of $V(G)$ with $A \cup B = V(G)$, such that no edge of G joins a vertex in $A - B$ to a vertex in $B - A$. Its *order* is $|A \cap B|$. A well-known theorem of Lipton and Tarjan [7] asserts the following. (R^+ denotes the set of non-negative real numbers. If $w : V(G) \to R^+$ is a function and $X \subseteq V(G)$, we denote $\sum(w(v) : v \in X)$ by $w(X)$.)

Theorem 1.1 (Lipton-Tarjan [7]) *Let G be a planar graph with n vertices, and let $w : V(G) \to R^+$ be a function. Then there is a separation (A, B) of G of order $\leq 2\sqrt{2}n^{1/2}$, such that $w(A - B), w(B - A) \leq \frac{2}{3}w(V(G))$.*

Here we prove an extension of this theorem to non-planar graphs with a fixed excluded "minor". A graph H is a *minor* of a graph G if H can be obtained from a subgraph of G by contracting edges. By an *H-minor* of G we mean a minor of G isomorphic to H. Thus, the Kuratowski-Wagner Theorem (see, e.g., [1]) asserts that planar graphs are those without K_5 or $K_{3,3}$ minors. We prove the following result.

Theorem 1.2 *Let $h \geq 1$ be an integer, let H be a simple graph with h vertices, and let G be a graph with n vertices, no H-minor and with a weight function $w : V(G) \to R^+$. Then there is a separation (A, B) of G of order $\leq h^{3/2}n^{1/2}$, such that $w(A - B), w(B - A) \leq \frac{2}{3}w(V(G))$.*

0

Note that since the complete graph K_h contains every simple graph on h vertices it suffices to prove this theorem for the case $H = K_h$. Hence we deal, from now on, only with this case.

We suspect that the estimate $h^{3/2}n^{1/2}$ in Theorem 1.2 can be replaced by $O(hn^{1/2})$. Since the genus of K_h is $\Theta(h^2)$ this would extend, if true, a result of Gilbert, Hutchinson and Tarjan [3] who proved that any graph on n vertices with genus g has a separator of order $O(g^{1/2}n^{1/2})$. Although an $O(hn^{1/2})$ result would be an asymptotically better result than the one stated above, an advantage of the formulation given in Theorem 1.2 is that there are no hidden constants in the expression for the size of the separator given by it.

If G is a graph and $X \subseteq V(G)$, an X-flap is the vertex set of some component of $G \setminus X$ (=the graph obtained from G by deleting X.) Let $w : V(G) \to R^+$ be a function. If $X \subseteq V(G)$ is such that $w(F) \leq \frac{2}{3}w(V(G))$ for every X-flap F then it is easy to find a separation (A, B) with $A \cap B = X$ such that $w(A - B), w(B - A) \leq \frac{2}{3}w(V(G))$. Thus, Theorem 1.2 is implied by the following.

Proposition 1.3 *Let $h \geq 1$ be an integer, let G be a graph with n vertices and with no K_h-minor and let $w : V(G) \to R^+$ be a weight function. Then there exists $X \subseteq V(G)$ with $|X| \leq h^{3/2}n^{1/2}$ such that $w(F) \leq \frac{1}{2}w(V(G))$ for every X-flap F.*

Our proof is algorithmic and in fact we can show:

Theorem 1.4 *There is an algorithm with running time $O(h^{1/2}n^{1/2}m)$, which takes as input an integer $h \geq 1$, a graph G (where $n = |V(G)|$ and $m = |V(G)| + |E(G)|$), and a function $w : V(G) \to R^+$. It outputs either*
(a) a K_h-minor of G, or
(b) a subset $X \subseteq V(G)$ with $|X| \leq h^{3/2}n^{1/2}$ such that $w(F) \leq \frac{1}{2}w(V(G))$ for every X-flap F.

We note that (unlike some other recent polynomial-time algorithms involving graph minors) this algorithm is practical and easy to implement, and there are no large constants hidden in the O notation above.

Note that our algorithm is not as efficient as the linear time one given by Lipton and Tarjan for the planar case. In fact, even for the case of bounded (orientable) genus, the algorithm of [3] for finding a small separator runs in linear time, given an embedding of the graph in its genus surface. (It is not known, though, how to find such an embedding in time which is polynomial in both the genus and the size of the graph, and in fact this is impossible if $P \neq NP$, since the problem of determining the genus of a graph is, as proved by C. Thomassen [11], NP-complete. Moreover, even for bounded genus, the best known algorithm for finding such an embedding is much slower than our separator algorithm). Although our algorithm is not as efficient, it is much more general. As a very special case it gives an algorithm for finding small separators in graphs with genus g, which is polynomial in both n and g, even when we are not given an embedding of the graph in its genus surface. The question of finding such an algorithm was raised in [3].

By a *haven of order k* in G we mean a function β which assigns to each subset $X \subseteq V(G)$ with $|X| \leq k$ an X-flap $\beta(X)$, in such a way that if $X \subseteq Y$ and $|Y| \leq k$ then $\beta(Y) \subseteq \beta(X)$. Clearly, if Proposition 1.3 is false for G, then for each $X \subseteq V(G)$ with $|X| \leq h^{3/2}n^{1/2}$ there is a unique X-flap, say $\beta(X)$, with $w(\beta(X)) > \frac{1}{2}w(V(G))$; and β thus defined is evidently a haven of order $h^{3/2}n^{1/2}$. Therefore, Proposition 1.3 is implied by the following more general and more compact result.

Theorem 1.5 *Let $h \geq 1$ be an integer, and let G be a graph with n vertices with a haven of order $h^{3/2}n^{1/2}$. Then G has a K_h-minor.*

Lipton and Tarjan [8] , [9] and Lipton, Rose and Tarjan [6] gave many applications of the planar separator theorem (and noted that most of them would generalize to any family of graphs with small separators.) Indeed our results supply simple generalizations of all these applications. In particular it follows that for any fixed graph H , given a graph G with n vertices and with no H-minor one can approximate the size of the maximum independent set of G up to a relative error of $1/\sqrt{\log n}$ in polynomial time. In time $2^{O(\sqrt{n})}$ one can find that size exactly and find the chromatic number of G. Also, any sparse system of n linear equations in n unknowns whose sparsity structure corresponds to G can be solved in time $O(n^{3/2})$. It can also be shown that graphs with an excluded minor are easy to pebble, can be imbedded with a small average proximity in binary trees, and cannot be the underlying graphs of efficient Boolean circuits computing certain functions.

The rest of this paper is organized as follows. In the next two sections we present the proofs of the main results; Theorem 1.4 and Theorem 1.5. Our method, unlike the ones used for proving the separator theorems for planar graphs and for graphs of bounded genus, does not involve any topological considerations concerning the embedding of the graphs, and is purely combinatorial. A simple but useful tool is a lemma concerning connecting trees presented in Section 2. In Section 4 we discuss the applications of our results; outline the extensions of the applications given in [9] and [6] to graphs with excluded minors and mention another combinatorial application that relates the tree-width of a graph and the existence of minors of complete graphs in it.

2 Finding small connecting trees

We shall need the following lemma.

Lemma 2.1 *Let G be a graph with n vertices, let $A_1, \ldots, A_k$ be k subsets of $V(G)$ and let $r \geq 1$ be a real number. Then either*
(i) there is a tree T in G with $|V(T)| \leq r$ such that $V(T) \cap A_i \neq \emptyset$ for $i = 1, \ldots, k$ or
(ii) there exists $Z \subseteq V(G)$ with $|Z| \leq (k-1)n/r$, such that no Z-flap intersects all of $A_1, \ldots, A_k$.

Proof.   We may assume that $k \geq 2$. Let $G^1, \ldots, G^{k-1}$ be isomorphic copies of G, mutually disjoint. For each $v \in V(G)$ and $1 \leq i \leq k-1$, let v^i be the corresponding vertex of G^i. Let J be the graph obtained from $G^1 \cup \ldots \cup G^{k-1}$ by adding, for $2 \leq i \leq k-1$ and all $v \in A_i$, an edge joining v^{i-1} and v^i. Let $X = \{v^1 : v \in A_1\}$, and $Y = \{v^{k-1} : v \in A_k\}$. For each $u \in V(J)$, let $d(u)$ be the number of vertices in the shortest path of J between X and u (or ∞ if there is no such path). There are two cases:
Case 1: $d(u) \leq r$ for some $u \in Y$.
Let P be a path of J between X and Y with $\leq r$ vertices. Let $S = \{v \in V(G) : v^i \in V(P) \; for \; some \; i, \; 1 \leq i \leq k-1\}$. Then $|S| \leq |V(P)| \leq r$, the subgraph of G induced on S is connected, and $S \cap A_i \neq \emptyset$ for $1 \leq i \leq k$. Thus (i) holds.
Case 2: $d(u) > r$ for all $u \in Y$.
Let t be the least integer with $t \geq r$. For $1 \leq j \leq t$, let $Z_j = \{u \in V(J) : d(u) = j\}$. Since $|V(J)| = (k-1)n$ and $Z_1, \ldots, Z_t$ are mutually disjoint, one of them, say Z_j, has cardinality $\leq (k-1)n/t \leq (k-1)n/r$. Now every path of J between X and

Y has a vertex in Z_j, for $d(u) \geq j$ for all $u \in Y$. Let $Z = \{v \in V(G) : v^i \in Z_j \; for \; some \; i \, , \; 1 \leq i \leq k-1\}$. Then $|Z| \leq |Z_j| \leq (k-1)n/r$, and we claim that Z satisfies (ii). For suppose that F is a Z-flap of G which intersects all of $A_1, \ldots, A_k$. Let $a_i \in F \cap A_i$ $(1 \leq i \leq k)$, and for $1 \leq i \leq k-1$ let P_i be a path of G with $V(P_i) \subseteq F$ and with ends a_i, a_{i+1}. Let P^i be the path of G^i corresponding to P_i. Then $V(P^1) \cup \ldots \cup V(P^{k-1})$ includes the vertex set of a path of J between X and Y, and yet is disjoint from Z_j, a contradiction. Thus, there is no such F, and so (ii) holds. □

We observe that the last proof is easily converted to an algorithm with running time $O(km)$, which, with input G, r and $A_1, \ldots, A_k$ as in the lemma (where $m = |V(G)| + |E(G)|$), computes either a tree T as in (i) or a set Z as in (ii).

3 The algorithm and the proof

First, we present the proof of Theorem 1.5, and then adapt the proof to yield an algorithm for Theorem 1.4. Let G be a graph. By a *covey* in G we mean a set $\mathcal{C}$ of (non-null) trees of G, mutually vertex-disjoint, such that for all distinct $C_1, C_2 \in \mathcal{C}$ there is an edge of G with one end in $V(C_1)$ and the other in $V(C_2)$. Thus, if G has a covey of cardinality h then it has a K_h-minor.
Proof of Theorem 1.5 Let β be a haven in G of order $h^{3/2}n^{1/2}$. Choose $X \subseteq V(G)$ and a covey $\mathcal{C}$ with $|\mathcal{C}| \leq h$ such that
(i) $X \subseteq \cup(V(C) : C \in \mathcal{C})$
(ii) $|X \cap V(C)| \leq h^{1/2}n^{1/2}$ for each $C \in \mathcal{C}$
(iii) $V(C) \cap \beta(X) = \emptyset$ for each $C \in \mathcal{C}$
(iv) subject to (i), (ii) and (iii), $|\mathcal{C}| + 2(|\beta(X)| + |X \cup \beta(X)|)$ is minimum.
     (This is certainly possible, for $\mathcal{C} = X = \emptyset$ satisfy (i), (ii) and (iii).) Let $\mathcal{C} = \{C_1, \ldots, C_k\}$. We suppose for a contradiction that $k < h$. For $1 \leq i \leq k$, let A_i be the set of all $v \in \beta(X)$ adjacent in G to a vertex of C_i. Let G' be the restriction of G to $\beta(X)$. By Lemma 2.1 applied to G' with $r = h^{1/2}n^{1/2}$, one of the following cases holds:
Case 1: there is a tree T of G' with $|V(T)| \leq h^{1/2}n^{1/2}$, such that $V(T) \cap A_i \neq \emptyset$ for $1 \leq i \leq k$.
Case 2: there exists $Z \subseteq \beta(X)$ with $|Z| \leq (k-1)|\beta(X)|/h^{1/2}n^{1/2} \leq h^{1/2}n^{1/2}$ such that no Z-flap of G' intersects all of $A_1, \ldots, A_k$.
     In the first case we replace $\mathcal{C}$ by $\mathcal{C}' = \mathcal{C} \cup \{T\}$ and X by $X' = X \cup V(T)$. It can be easily checked

that this supplies a contradiction to (iv).

The second case is somewhat more complicated. In this case we can enlarge one of the members of $\mathcal{C}$ to include some vertices of Z together with several additional properly chosen vertices. This is done as follows. Define $Y = X \cup Z$. Since $k \leq h - 1$ it follows that $|Y| \leq h^{3/2}n^{1/2}$, and so $\beta(Y)$ exists and $\beta(Y) \subseteq \beta(X)$. Since $\beta(Y)$ is a Z-flap of G' there exists an index i with $1 \leq i \leq k$ such that $\beta(Y) \cap A_i = \emptyset$. Extend C_i to a maximal tree C_i' of G disjoint from $\beta(Y)$ and from each C_j, $(j \neq i)$. Let $Z' = V(C_i') \cap Z$, let $X' = Z' \cup (X - V(C_i))$, and let $W = V(C_i') \cup (V(G) - \beta(X))$.

We claim that $\beta(X') \cap W = \emptyset$. For suppose not. Since $\beta(Y) \subseteq \beta(X')$, there is a path of G between W and $\beta(Y)$ contained within $\beta(X')$ and hence disjoint from X'. Since $W \cap \beta(Y) = \emptyset$, there are two consecutive vertices u,v of this path with $u \in W$ and $v \in V(G) - W \subseteq \beta(X)$. Since u,v are adjacent it follows that $u \in X \cup \beta(X)$, and so

$$u \in (X \cup \beta(X)) \cap (W - X') \subseteq V(C_i').$$

Since $v \notin W$ it follows from the maximality of C_i' that $v \in \beta(Y)$. Since $u \notin \beta(Y)$ we deduce that $u \in Y$, and so

$$u \in Y \cap (V(C_i') - X') \subseteq V(C_i).$$

But then $v \in A_i$, which is impossible, since $A_i \cap \beta(Y) = \emptyset$. This proves our claim that $\beta(X') \cap W = \emptyset$. Hence $\beta(X') \subseteq \beta(X)$. Let $\mathcal{C}' = (\mathcal{C} - \{C_i\}) \cup \{C_i'\}$; then $\mathcal{C}'$ is a covey. We observe that
(i) $X' \subseteq \cup(V(C) : C \in \mathcal{C}')$; for $Z' \subseteq V(C_i')$
(ii) $|X' \cap V(C)| \leq h^{1/2}n^{1/2}$ for each $C \in \mathcal{C}'$; for if $C \neq C_i'$ then $X' \cap V(C) = X \cap V(C)$, and $X' \cap V(C_i') = Z'$
(iii) $V(C) \cap \beta(X') = \emptyset$ for each $C \in \mathcal{C}'$; for $\beta(X') \cap W = \emptyset$, as we have seen.
By (iv),

$$|\mathcal{C}'| + 2(|\beta(X')| + |X' \cup \beta(X')|) \geq$$

$$|\mathcal{C}| + 2(|\beta(X)| + |X \cup \beta(X)|).$$

But $|\mathcal{C}'| = |\mathcal{C}|$, $\beta(X') \subseteq \beta(X)$ and $X' \cup \beta(X') \subseteq (X \cup \beta(X)) - (X \cap V(C_i))$, and so $X \cap V(C_i) = \emptyset$. Then $\mathcal{C} - \{C_i\}, X$ satisfy (i),(ii) and (iii), contrary to (iv).

In both cases, therefore, we have obtained a contradiction. Thus our assumption that $k < h$ was incorrect, and so $k = h$ and G has a K_h-minor, as required. $\square$

Now let us convert the proof above to an algorithm, as promised in Theorem 1.4. Let h, G, w be the input. Set $X_0 = \mathcal{C}_0 = \emptyset$ and $B_0 = V(G)$, and begin the first iteration. In general, at the beginning of the tth iteration, we have a subset $X_{t-1} \subseteq V(G)$, a covey $\mathcal{C}_{t-1}$ with $|\mathcal{C}_{t-1}| \leq h$, and a subset $B_{t-1} \subseteq V(G)$ which is a union of X_{t-1}-flaps, such that
(i) $X_{t-1} \subseteq \cup(V(C) : C \in \mathcal{C}_{t-1})$
(ii) $|X_{t-1} \cap V(C)| = [h^{1/2}n^{1/2}]$ for each $C \in \mathcal{C}_{t-1}$
(iii) $V(C) \cap B_{t-1} = \emptyset$ for each $C \in \mathcal{C}_{t-1}$
(iv) $w(F) \leq \frac{1}{2}w(V(G))$ for each X_{t-1}-flap F which is not a subset of B_{t-1}.

1. Let $|\mathcal{C}_{t-1}| = k$. If $k = h$ we have found a K_h-minor; we output (a) and stop. Otherwise we go to step (2).

2. Compute the X_{t-1}-flaps of the induced subgraph of G on B_{t-1}. If $w(F) \leq \frac{1}{2}w(V(G))$ for each such flap F we output (b) (with $X = X_{t-1}$) and stop. Otherwise we let F be the unique X_{t-1}-flap with $w(F) > \frac{1}{2}w(V(G))$ and go to step (3).

3. Let $\mathcal{C}_{t-1} = \{C_1, \ldots, C_k\}$. If $|F| \leq [h^{1/2}n^{1/2}]$ we output (b), with $X = X_{t-1} \cup F$, and stop. Otherwise we let G' be the induced subgraph of G on F. For $1 \leq i \leq k$, let $A_i \subseteq V(G')$ be the set of all $v \in V(G')$ with a neighbour in $V(C_i)$. If $A_i = \emptyset$ for some i, we set $X_t = X_{t-1} - V(C_i)$, $\mathcal{C}_t = \mathcal{C}_{t-1} - \{C_i\}$ and $B_t = F$, and return to step (1) for the next iteration. Otherwise, we go to step (4).

4. We apply Lemma 2.1 to G' and $A_1, \ldots, A_k$, taking $r = h^{1/2}n^{1/2}$. We obtain either:
(i) a tree T of G' with $|V(T)| \leq h^{1/2}n^{1/2}$ such that $V(T) \cap A_i \neq \emptyset$ for each i, or
(ii) a subset $Z \subseteq V(G')$ such that no Z-flap of G' intersects all of $A_1, \ldots, A_k$.
In the first case we go to step (5), and in the second to step (6).

5. Given T as in (4)(i), extend it to a tree T' of size $[h^{1/2}n^{1/2}]$ in G' and then set $X_t = X_{t-1} \cup V(T')$, $B_t = F - V(T')$, and $\mathcal{C}_t = \mathcal{C}_{t-1} \cup \{T'\}$, and return to step (1) for the next iteration.

6. Given Z as in (4)(ii), let $Y = X_{t-1} \cup Z$. If $w(L) \leq \frac{1}{2}w(V(G))$ for every Y-flap L, we output (b) (with $X = Y$) and stop. Otherwise, there is a unique Y-flap L with $w(L) >$

$\frac{1}{2}w(V(G))$, and it satisfies $L \subseteq B_{t-1}$. Choose i with $1 \leq i \leq k$ such that $L \cap A_i = \emptyset$. Extend C_i to a maximal tree C_i' of G with the property that $V(C_i')$ is disjoint from L and from each $V(C_j)$ ($j \neq i$). Put $Z' = Z \cap V(C_i')$. If $Z' < [h^{1/2}n^{1/2}]$ extend it to a set of that size by adding to it sufficiently many vertices from F such that the graph induced on $V(C_i') \cup Z'$ would still be connected and replace C_i' by a spanning tree on this union. Let $X_t = (X_{t-1} - V(C_i)) \cup Z'$, and $\mathcal{C}_t = (\mathcal{C}_{t-1} - \{C_i\}) \cup \{C_i'\}$, and let B_t be the X_t-flap of maximum weight among those contained in F. We return to step (1) for the next iteration.

This completes the description of the algorithm. Its correctness can be argued as in the proof of Theorem 1.5. We have omitted details of the (obvious) data structures used, but if they are chosen appropriately, then each of the steps (1), (2), (3), (5) and (6) takes time $O(m)$ in each iteration in which it is called, and step (4) takes time $O(hm)$ in each iteration it is called. Notice that in each iteration the quantity $|B_t| + |B_t \cup X_t|$ decreases by at least $[h^{1/2}n^{1/2}]$ (and this quantity never increases during the algorithm). It follows that each of the steps (1)-(6) is performed $O(h^{-1/2}n^{1/2})$ times. The algorithm thus has running time $O(h^{1/2}n^{1/2}m)$, as required. This completes the proof of Theorem 1.4. $\square$

It may be that by using more sophisticated, dynamic data structures, a more efficient implementation of the algorithm can be found. At the moment this remains an open problem.

4   Applications

Efficient algorithms for finding small separators in graphs are useful in the layout of circuits in a model of VLSI (see, e.g., [5]). Thus our results can be applied for finding efficiently embeddings of graphs with excluded minors.

All the applications of the Lipton-Tarjan planar separator theorem given in [9] and in [6] carry over, by our result, to any class of graphs with an excluded minor. Since most of the proofs are straightforward extensions of those given in the above papers we merely state each result, and only add a few words about its proof when it is not an obvious

extension of the one for the planar case. In all the propositions in this section, when G is a graph we let m denote the sum of the number of its vertices and the number of its edges.

By repeatedly applying the separating algorithm of Theorem 1.5 to a graph one can obtain the following immediate generalization of a result of [9].

Proposition 4.1 *Let G be an n-vertex graph with no K_h-minor, and with nonnegative weights whose total sum is 1 assigned to its vertices. Then, for any $0 < \epsilon \leq 1$ there is a set of at most $O(h^{3/2}n^{1/2}/\epsilon^{1/2})$ vertices of G whose removal leaves G with no connected component whose total weight exceeds ϵ. Such a set can be found in time $O(h^{1/2}n^{1/2}m)$, (independent of ϵ).*

This proposition can be used to obtain a polynomial time algorithm for approximating the size of the maximum independent set of a graph with an excluded minor, (and for approximating several similar quantities, whose exact determination is known to be NP-complete). We simply break the graph into pieces of small size (say, size $\frac{1}{2}\log n$), find the maximum independent set in each piece by exhaustive search, and combine the results to obtain the desired approximation. To estimate the relative error here one can apply the result of Kostochka [4] and Thomason [12] and a simple greedy-argument to obtain a lower bound for the size of the maximum independent set in an n-vertex graph with no K_h-minor. This gives:

Proposition 4.2 *There is an algorithm that approximates, given an n-vertex graph G with no K_h-minor , the size of a maximum independent set in it with a relative error of $O(\frac{h^{5/2}(\log h)^{1/2}}{(\log n)^{1/2}})$ in time $O(h^{1/2}n^{1/2}m)$.*

Separator theorems are useful in designing efficient divide-and-conquer algorithms. An example given in [9] is that of nonserial dynamic programming (see, e.g., [13]).

Proposition 4.3 *Let $f(x_1, \ldots, x_n)$ be a function of n variables, where each variable takes values in a finite set S of s elements, and suppose f is a sum of functions f_i, where each f_i is a function of some of the variables x_j. Let G be the graph whose vertices are the variables x_j in which x_j is adjacent to x_k iff they both appear in a common term f_i. Suppose , further, that G has no K_h-minor. Then one can*

find the maximum of f in the domain $x_j \in S$ for all j in time $s^{O(h^{3/2}n^{1/2})}$.

Note that the graph G is not planar if a single term f_i is a function of more than 4 variables, whereas the extension to graphs with an excluded minor allows much more complicated functions f_i.

It is not too difficult to see that the last proposition implies that one can find the maximum size of an independent set in an n-vertex graph G with no K_h-minor in time $2^{O(h^{3/2}n^{1/2})}$ and check if such a graph is s-colorable in time $s^{O(h^{3/2}n^{1/2})}$. For fixed h this shows that the maximum size of an independent set in a graph with no K_h-minor and with $\log^2 n$ vertices can be found in time $n^{O(1)}$. This can be applied to modify the proof of Proposition 4.2 and obtain a polynomial-time approximation algorithm for the size of the maximum independent set in an n-vertex graph with a fixed excluded minor with relative error of $O(\frac{1}{\log n})$.

Pebbling is a one person game that arises in the study of time-space trade-offs (see, e.g., [2]). An immediate extension of a result from [9] gives;

Proposition 4.4 *Any n-vertex acyclic digraph with no K_h-minor and with maximum in-degree k can be pebbled using $O(h^{3/2}n^{1/2} + k \log n)$ pebbles.*

Other straight-forward extensions of known applications are the fact that if h is fixed and G has no K_h-minor then:
(i) Any Boolean circuit whose underlying graph is G which computes the product of two m-bit integers has at least $\Omega(m^2)$ vertices.
(ii)If k is the maximum degree of a vertex in G then G can be embedded in a binary tree with $O(k)$ average-proximity, i.e., there is a one-to-one mapping of the set of vertices of G into that of a binary tree such that the average distance (in the tree) between the images of two neighbours in G is $O(k)$.

The result of Lipton, Rose and Tarjan [6] also generalizes easily; one can solve any sparse system of n linear equations in n unknowns whose sparsity structure corresponds to a graph on n vertices with a fixed excluded minor in time $O(n^{3/2})$. (The graph here is the one whose vertices are $1, \ldots, n$ and i is adjacent to j iff the coefficient of the ith variable in the jth equation is non-zero (the system is symmetric, and hence this is a well-defined graph)). The algorithm suggested by Gabow (cf. [9]) for finding a maximum matching also extends to the fixed-excluded-minor case; although its running time matches that of the best general algorithm it seems simpler to implement.

Finally we mention a combinatorial application of our results. A *tree-decomposition* of a graph G is a pair (T, W), where T is a tree and $W = (W_t : t \in V(T))$ is a family of subsets of $V(G)$, such that $\cup(W_t : t \in V(T)) = V(G)$, for every $e \in E(G)$ there exists $t \in V(T)$ such that W_t contains both ends of e, and if $t_1, t_2, t_3 \in V(T)$ and t_2 lies on the path between t_1 and t_3 then $W_{t_1} \cap W_{t_3} \subseteq W_{t_2}$. The *tree-width* of G is the minimum k such that there is a tree-decomposition (T, W) of T satisfying $|W_t| \leq k + 1$ for all $t \in V(T)$. Combining Theorem 1.5 with one of the results of [10] we can prove the following.

Proposition 4.5 *Let $h \geq 1$ be an integer, and let G be a graph with n vertices and with tree-width at least $h^{3/2}n^{1/2}$. Then G has a K_h-minor.*

References

[1] J. A. Bondy and U. S. R. Murty, *Graph Theory with Applications*, The Macmillan Press, London and Basingstoke, 1978.

[2] S. A. Cook, *An observation on time-storage trade-offs* , Proc. Fifth Annual STOC (1973), 29-33.

[3] J. R. Gilbert, J. P. Hutchinson and R. E. Tarjan, *A separator theorem for graphs of bounded genus*, J. Algorithms 5(1984), 391-407.

[4] A. V. Kostochka, *A lower bound for the Hadwiger number of a graph as a function of the average degree of its vertices*, Diskret. Analiz. Novosibirsk 38(1982), 37-58. (In Russian).

[5] C. E. Leiserson, *Area efficient graph layouts (for VLSI)*, Proc. 21 FOCS (1980), 270-281.

[6] R. J. Lipton, D. J. Rose and R. E. Tarjan, *Generalized nested dissection*, SIAM J. Numer. Anal., 16(1979), 177-189.

[7] R. J. Lipton and R. E. Tarjan, *A separator theorem for planar graphs*, SIAM J. Appl. Math. 36(1979), 177-189.

[8] R. J. Lipton and R. E. Tarjan, *Applications of a planar separator theorem*, Proc. 18th FOCS (1977), 162-170.

[9] R. J. Lipton and R. E. Tarjan, *Applications of a planar separator theorem*, SIAM J. Comput. 9(1980), 615-627.

[10] P. D. Seymour and R. Thomas, *Graph searching, and a minimax theorem for tree-width*, to appear.

[11] C. Thomassen, *The graph genus problem is NP-complete*, J. of Algorithms 10(1989), 568-576.

[12] A. G. Thomason, *An extremal function for contractions of graphs*, Math. Proc. Cambridge Philos. Soc. 95(1984), 261-265.

[13] A. Rosenthal, *Nonserial dynamic programming is optimal*, Proc. Ninth Annual ACM STOC (1977), 98-105.