

A Linear-Time Algorithm to Find a Separator in a Graph Excluding a Minor

BRUCE REED

McGill University and Centre National de la Recherche Scientifique

AND

DAVID R. WOOD

The University of Melbourne

Abstract. Let G be an n -vertex m -edge graph with weighted vertices. A pair of vertex sets $A, B \subseteq V(G)$ is a $\frac{2}{3}$ -separation of order $|A \cap B|$ if $A \cup B = V(G)$, there is no edge between $A - B$ and $B - A$, and both $A - B$ and $B - A$ have weight at most $\frac{2}{3}$ the total weight of G . Let $\ell \in \mathbb{Z}^+$ be fixed. Alon et al. [1990] presented an algorithm that in $\mathcal{O}(n^{1/2}m)$ time, outputs either a K_ℓ -minor of G , or a separation of G of order $\mathcal{O}(n^{1/2})$. Whether there is a $\mathcal{O}(n + m)$ -time algorithm for this theorem was left as an open problem. In this article, we obtain a $\mathcal{O}(n + m)$ -time algorithm at the expense of a $\mathcal{O}(n^{2/3})$ separator. Moreover, our algorithm exhibits a trade-off between time complexity and the order of the separator. In particular, for any given $\epsilon \in [0, \frac{1}{2}]$, our algorithm outputs either a K_ℓ -minor of G , or a separation of G with order $\mathcal{O}(n^{(2-\epsilon)/3})$ in $\mathcal{O}(n^{1+\epsilon} + m)$ time. As an application we give a fast approximation algorithm for finding an independent set in a graph with no K_ℓ -minor.

Categories and Subject Descriptors: G.2.2 [Discrete Mathematics]: Graph Theory—Graph algorithms; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—Computations on discrete structures

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Graph, minor, separation, separator

39

B. Reed is the Canada Research Chair in Graph Theory, supported by the Natural Sciences and Engineering Research Council of Canada. D. Wood is a QEII Research Fellow, supported by the Australian Research Council.

A preliminary version of this article was published in *Proceedings of the 2005 European Conference on Combinatorics, Graph Theory and Applications (EuroComb'05)*, Discrete Mathematics and Theoretical Computer Science, vol. AE, 45-50.

Authors' addresses: B. Reed, School of Computer Science, McGill University, Montréal, Canada and Laboratoire I3S, Centre National de la Recherche Scientifique, Sophia-Antipolis, France, e-mail: breed@cs.mcgill.ca; D. Wood, Department of Mathematics and Statistics, The University of Melbourne, Melbourne, Australia, e-mail: woodd@unimelb.edu.au.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2009 ACM 1549-6325/2009/10-ART39 \$10.00

DOI 10.1145/1597036.1597043 <http://doi.acm.org/10.1145/1597036.1597043>

ACM Reference Format:

Reed, B. and Wood, D. R. 2009. A linear-time algorithm to find a separator in a graph excluding a minor. *ACM Trans. Algor.* 5, 4, Article 39 (October 2009), 16 pages.
 DOI = 10.1145/1597036.1597043 <http://doi.acm.org/10.1145/1597036.1597043>

1. Introduction

This article presents a linear-time algorithm for finding a separator in graphs excluding a fixed minor.


A *separation* of a graph¹ G is a pair $\{A, B\}$ of vertex sets $A, B \subseteq V(G)$ such that $A \cup B = V(G)$, and there is no edge between $A - B$ and $B - A$, as illustrated in Figure 1. The *order* of $\{A, B\}$ is $|A \cap B|$. The set $A \cap B$ is called a *separator* of G . A *weighting* of G is a function $w : V(G) \rightarrow \mathbb{R}^+$. Let $w(S) := \sum_{v \in S} w(v)$ for all $S \subseteq V(G)$, and let $w(G) := w(V(G))$. We say (G, w) is a *weighted graph*. A separation $\{A, B\}$ of a weighted graph (G, w) is a β -*separation* if $w(A - B) \leq \beta \cdot w(G)$ and $w(B - A) \leq \beta \cdot w(G)$.

A “separator theorem” is of the format: For some $0 < \beta < 1$ and $0 < \epsilon \leq 1$, every graph G from a certain family has a β -separation of order $\mathcal{O}(|G|^{1-\epsilon})$. Applications of separator theorems are numerous, and include VLSI circuit layout [Leiserson 1980], approximation algorithms using the divide-and-conquer paradigm [Chiba et al. 1981; Lipton and Tarjan 1980], solving sparse systems of linear equations [Lipton et al. 1979], pebbling games [Lipton and Tarjan 1980], and graph drawing [Dujmović and Wood 2004]. See the monograph by Rosenberg and Heath [2001] for more details.

A seminal theorem due to Lipton and Tarjan [1979] states that every weighted planar graph G has a $\frac{2}{3}$ -separation of order $\mathcal{O}(|G|^{1/2})$ that can be computed in $\mathcal{O}(|G| + \|G\|)$ time. The importance of this result cannot be overstated, as suggested by the amount of effort that has gone into improving the constant in the $\mathcal{O}(|G|^{1/2})$ bound [Chung 1991; Djidjev 1982; Alon et al. 1994; Venkatesan 1987; Djidjev 1987]. Many other aspects of separators in planar graphs have been studied. For example, Miller [1986] proved that every 2-connected planar graph has a cycle separator, and Djidjev and Venkatesan [1997] improved the constants. Aleksandrov et al. [2006] and Djidjev [2000] considered separators in planar graphs whose order is measured in terms of associated vertex costs.

Djidjev and Gilbert [1999] considered separators in graphs with negative and multiple weights. Separators in certain geometric graphs have been studied by Miller et al. [1997] and Smith and Wormald [1998]. Plaisted [1990] developed a heuristic for finding separators in arbitrary graphs. Edge separators have been studied by Sýkora and Vřto [1993] and Diks et al. [1993]. Alber et al. [2003] studied separators from the perspective of the theory of fixed parameter tractability. Approximation algorithms for separators are also well studied [Garg et al. 1999; Feige and Mahdian 2006; Arora et al. 2004; Amir et al. 2003; Even et al. 2000; Even et al. 1999; Bodlaender et al. 1995].

¹We consider graphs G that are simple, finite, and undirected. Let $V(G)$ and $E(G)$ denote the vertex and edge sets of G . Let $|G| := |V(G)|$ and $\|G\| := |E(G)|$. For a set $S \subseteq V(G)$, let $G[S]$ denote the subgraph of G induced by S . For each vertex $v \in V(G)$, let $N(v) := \{w \in V(G) : vw \in E(G)\}$ be the set of neighbors of v . For each subgraph X of G , let $N(X) := \bigcup\{N(v) - V(X) : v \in V(X)\}$. For $n \in \mathbb{Z}^+$, let $[n] := \{1, 2, \dots, n\}$.

FIG. 1. A separation $\{A, B\}$.

The theorem of Lipton and Tarjan was generalized to graphs with genus γ by Gilbert et al. [1984] and Djidjev [1987, 1985b, 1981]. They proved that such graphs G have a separation of order $\mathcal{O}(\gamma^{1/2} \cdot |G|^{1/2})$, which can be computed in linear time [Djidjev 1985a; Aleksandrov and Djidjev 1996]. The special case of toroidal graphs was considered by Aleksandrov and Djidjev [1989].

Perhaps the most general setting for separator theorems is for graphs excluding a fixed minor, as studied by Alon et al. [1990b], Plotkin et al. [1994], Grohe [2003], and Demaine and Hajiaghayi [2008a, 2008b, 2005]. A graph H is a *minor* of a graph G if a graph isomorphic to H can be obtained from a subgraph of G by contracting edges, in which case we say that G contains an H -minor. An H -model in G is a set of disjoint connected subgraphs $\{X_v : v \in V(H)\}$ indexed by the vertices of H , such that for every edge $vw \in E(H)$, there is an edge $xy \in E(G)$ with $x \in X_v$ and $y \in X_w$. Clearly G contains an H -minor if and only if G contains an H -model. For algorithmic purposes, we choose to work with H -models rather than H -minors. Graph classes defined by an excluded minor are often of interest. For example, the Kuratowski-Wagner theorem states that a graph is planar if and only if it contains no K_5 -minor and no $K_{3,3}$ -minor. Alon et al. [1990b] proved the following generalization of the Lipton-Tarjan separator theorem for graphs excluding an arbitrary minor.

THEOREM 1.1. [ALON ET AL. 1990B]. *There is an algorithm that, given $\ell \in \mathbb{Z}^+$ and a weighted graph (G, w) , outputs either:*

- (a) a K_ℓ -model of G , or
- (b) a $\frac{2}{3}$ -separation of (G, w) of order at most $\ell^{3/2} \cdot |G|^{1/2}$ in time $\mathcal{O}((\ell \cdot |G|)^{1/2} \cdot (|G| + \|G\|))$.

Suppose that ℓ is fixed. It follows from a result of Mader [1967] that Theorem 1.1 can be implemented in $\mathcal{O}(|G|^{3/2} + \|G\|)$ time; see Theorem 2.3. Alon et al. [1990b] left as an open problem whether linear $\mathcal{O}(|G| + \|G\|)$ time is possible. The main result of this article is the following partial answer to this question. We obtain linear time complexity at the expense of a slightly larger separator (and larger dependence on ℓ). Moreover, our algorithm exhibits a trade-off between time complexity (ranging from $\mathcal{O}(n)$ to $\mathcal{O}(n^{3/2})$) and the order of the separator (ranging from $\mathcal{O}(n^{2/3})$ to $\mathcal{O}(n^{1/2})$).

THEOREM 1.2. *There is an algorithm that, given $\epsilon \in [0, \frac{1}{2}]$, $\ell \in \mathbb{Z}^+$, and a weighted graph (G, w) , outputs either:*

- (a) a K_ℓ -model of G , or
 (b) a $\frac{2}{3}$ -separation of (G, w) of order at most $\ell^{3/2} \cdot 2^{(\ell^2+4)/2} \cdot |G|^{(2-\epsilon)/3}$
 in time $\mathcal{O}(\ell \cdot 2^{(3\ell^2+2\ell+6)/2} \cdot |G|^{1+\epsilon} + \ell \cdot \|G\|)$.

Note that for applications to divide-and-conquer algorithms a separation of order $\mathcal{O}(|G|^{1-\epsilon})$, for some constant $\epsilon > 0$, is all that is needed. For example, in Section 5 we apply Theorem 1.2 to obtain an approximation algorithm for the maximum independent set problem on graphs excluding a fixed minor that runs in near-linear time and has diminishing relative error. (A set of vertices I in a graph is *independent* if no two vertices in I are adjacent.) Theorem 1.2 has also recently been applied by Tazari and Müller-Hannemann [2009] and Yuster [2008] to obtain improved shortest-paths algorithms for graphs excluding a fixed minor, and by Yuster and Zwick [2007] to obtain the fastest known algorithm for finding a maximum matching in a graph excluding a fixed minor.

We now outline the idea behind the proof of Theorem 1.2 for fixed ℓ and with $\epsilon = 0$. Suppose that in $\mathcal{O}(|G| + \|G\|)$ time, we can find a partition $\{S_1, S_2, \dots, S_{|G|^{2/3}}\}$ of $V(G)$, such that each S_i induces a connected subgraph of G with $\mathcal{O}(|G|^{1/3})$ vertices. Let H be the weighted graph obtained from G by contracting each subgraph $G[S_i]$ to a vertex v_i with weight $w(v_i) = w(S_i)$. Then apply Theorem 1.1 to H to obtain either a K_ℓ -model in H which defines a K_ℓ -model in G , or a $\frac{2}{3}$ -separation $\{A, B\}$ of H with order $\mathcal{O}(|H|^{1/2}) = \mathcal{O}(|G|^{1/3})$, in which case $\{\bigcup\{S_i : v_i \in A\}, \bigcup\{S_i : v_i \in B\}\}$ is a $\frac{2}{3}$ -separation of G with order $\mathcal{O}(|G|^{2/3})$. The time complexity is $\mathcal{O}(|H|^{3/2} + \|H\|) \subseteq \mathcal{O}(|G| + \|G\|)$.

The proof of Theorem 1.2 is actually a little different from this outline. In particular, the subgraphs $G[S_i]$ will not necessarily be connected. However, the partition of $V(G)$ will be “knitted” (see Section 4 for the definition), which will enable the output from Theorem 1.1 applied to H to be converted to the desired output for G . By relaxing the connectivity condition, we are able to prove that an appropriate partition exists.

In Section 2 we give an algorithmic version of a theorem of Mader [1967], which is used in Section 3 to prove an upper bound on the number of cliques in a graph excluding a minor. The main steps in the proof of Theorem 1.2 are presented in Section 4.

2. Mader’s Theorem

Mader [1967] proved that every sufficiently dense graph contains a large complete graph as a minor. In this section we prove the following algorithmic version of this result. Note that Robertson and Seymour [1995, page 85] proved a similar result with quadratic time complexity.

THEOREM 2.1. *Given a graph G with $\|G\| \geq 2^{\ell-3} \cdot |G|$ for some $\ell \in \mathbb{Z}^+$, a K_ℓ -model in G can be computed in $\mathcal{O}(\ell(|G| + \|G\|))$ time.*

Note that if we ignore the time complexity, Theorem 2.1 is far from best possible. Kostochka [1982, 1984] and Thomason [1984] independently proved that if $\|G\| \in \Omega(\ell\sqrt{\log \ell} \cdot |G|)$ then G contains a K_ℓ -model. In particular, Thomason [2001] proved that if $\|G\| \geq (\delta + o(1))\ell\sqrt{\log \ell} \cdot |G|$, where $\delta = 0.319\dots$ is a constant, then G contains a K_ℓ -model.




FIG. 2. Illustration of the proof of Lemma 2.2.

The proof of Theorem 2.1 is based on the following lemma.

LEMMA 2.2. *The following algorithm, given a graph G with $\|G\| \geq t \cdot |G|$ for some $t \in \mathbb{Z}^+$, outputs a connected nonempty induced subgraph X of G in time $\mathcal{O}(|G| + \|G\|)$, such that $G[N(X)]$ has minimum degree at least t .*

```

1: let  $U$  be a component of  $G$  with  $\|U\| \geq t \cdot |U|$ 
2: initialize  $X := G[\{v\}]$  for some vertex  $v \in V(U)$ 
3: while some vertex  $y \in N(X)$  has degree at most  $t - 1$  in  $G[N(X)]$  do
4:    $X := G[V(X) \cup \{y\}]$ 
5: end while
6: output  $X$ 

```

PROOF. To prove the correctness of the algorithm it suffices to show that, upon termination, $X \neq U$ and $N(X) \neq \emptyset$, implying that $G[N(X)]$ has minimum degree at least t . We do so, by showing that the invariant

$$e(X) \leq t(|X| - 1) + |N(X)| \quad (1)$$

is maintained, where $e(X)$ is the number of edges of U with at least one endpoint in X . Certainly (1) holds when $X = \{v\}$, in which case $e(X) = |N(X)| = \deg(v)$. Now suppose that (1) holds for some subgraph X of U , and $y \in N(X)$ has degree at most $t - 1$ in $G[N(X)]$. Let $X' := G[V(X) \cup \{y\}]$. Partition $N(y) - V(X)$ into two sets, $B := N(y) \cap N(X)$ and $C := N(y) - (V(X) \cup N(X))$, as illustrated in Figure 2. Since $|B| \leq t - 1$ and $N(X') = (N(X) - \{y\}) \cup C$,

$$\begin{aligned} e(X') &= e(X) + |B| + |C| \leq t(|X| - 1) + |N(X)| + t - 1 + |C| \\ &= t \cdot |X'| + |N(X')|. \end{aligned}$$

That is, (1) is satisfied for X' . Hence (1) is maintained throughout the algorithm. Now observe that $e(U) = \|U\| \geq t \cdot |U|$ and $N(U) = \emptyset$. Thus (1) is not satisfied for $X = U$. Hence, upon termination, $X \neq U$ and $N(X) \neq \emptyset$, and the algorithm computes X and $N(X)$ as claimed.

The algorithm can be implemented in $\mathcal{O}(|G| + \|G\|)$ time by maintaining the set $V(X)$, the set $N(X)$, the degree of each vertex in $G[N(X)]$, and a list L of the vertices in $N(X)$ with degree at most $t - 1$ in $G[N(X)]$. Whenever a vertex is moved from $N(X)$ into X or from $V(U) - (X \cup N(X))$ into $N(X)$, we traverse its list of neighbors, updating the degree within $N(X)$, and if necessary updating the list L . Thus, each list of neighbors is traversed $\mathcal{O}(1)$ times. Thus the algorithm can be implemented in $\mathcal{O}(|G| + \|G\|)$ time. We omit the routine description of the data structure manipulation necessary. \square

PROOF OF THEOREM 2.1. Theorem 2.1 is trivial for $\ell \leq 2$. Now assume that $\ell \geq 3$. Applying Lemma 2.2 with $t = 2^{\ell-3} (\geq 1)$, we obtain a nonempty connected subgraph X of G such that $G[N(X)]$ has minimum degree at least $2^{\ell-3}$. Thus $\|G[N(X)]\| \geq 2^{\ell-4}|N(X)|$. By induction, there is a $K_{\ell-1}$ -model in $G[N(X)]$. Since every vertex in $N(X)$ is adjacent to some vertex in X , this $K_{\ell-1}$ -model along with X forms a K_ℓ -model in G . There are ℓ applications of Lemma 2.2, each requiring $\mathcal{O}(|G| + \|G\|)$ time. \square

Theorem 2.1 implies the following slightly faster version of Theorem 1.1 (for fixed ℓ).

THEOREM 2.3. *There is an algorithm that, given $\ell \in \mathbb{Z}^+$ and a weighted graph (G, w) , outputs either:*

- (a) a K_ℓ -model of G , or
- (b) a $\frac{2}{3}$ -separation of (G, w) of order at most $\ell^{3/2} \cdot |G|^{1/2}$ in time $\mathcal{O}(\ell \cdot 2^\ell \cdot |G|^{3/2} + \ell \cdot \|G\|)$.

PROOF. If $\|G\| \geq 2^{\ell-3}|G|$, then a K_ℓ -model in G can be found in $\mathcal{O}(\ell(|G| + \|G\|))$ time by Theorem 2.1. Otherwise $\|G\| < 2^{\ell-3}|G|$, and the result follows from Theorem 1.1. \square

3. Cliques in Graphs Excluding a Minor

A critical aspect of the proof of our main result (Theorem 1.2) is an upper bound on the number of cliques in a graph excluding a given minor. We prove this bound in this section.

Let G be a graph. A k -clique of G is a (not necessarily maximal) set of k pairwise adjacent vertices of G . If every subgraph of G has a vertex of degree at most d , then G is d -degenerate. For example, Theorem 2.1 implies that a graph with no K_ℓ -minor is $2^{\ell-2}$ -degenerate.

We have the following crude bound on the number of cliques in a degenerate graph; see Wood [2007] and Norine et al. [2006] for similar results.

LEMMA 3.1. *A d -degenerate graph G with no k -clique has fewer than $d^{k-1} \cdot |G|$ cliques.*

PROOF. Since G is d -degenerate, we can order the vertices so that each vertex v has at most d neighbors to the left of v . Thus for all $i \in [k-1]$, every vertex is the rightmost vertex of at most $\binom{d}{i-1} \leq d^{i-1}$ cliques on i vertices. Thus every vertex is the rightmost vertex of at most $\sum_{i=1}^{k-1} d^{i-1} < d^{k-1}$ cliques. The result follows. \square

For example, a graph G with no K_ℓ -minor has fewer than $2^{(\ell-2)(\ell-1)} \cdot |G|$ cliques.

LEMMA 3.2. *Given a graph G with no k -clique and at least $2^{(\ell-2)(k-1)} \cdot |G|$ cliques for some $\ell, k \in \mathbb{Z}^+$, a K_ℓ -minor of G can be computed in $\mathcal{O}(\ell(|G| + \|G\|))$ time.*

PROOF. By Lemma 3.1 with $d = 2^{\ell-2}$, G is not $2^{\ell-2}$ -degenerate. By Lemma A.1 in Appendix A, a subgraph H of G with minimum degree greater than $2^{\ell-2}$ can be computed in $\mathcal{O}(|G| + \|G\|)$ time. Now $\|H\| > 2^{\ell-3} \cdot |H|$.




FIG. 3. A knitted C_4 -partition; each disc represents a connected component of a part of the partition.

Thus, by Theorem 2.1, a K_ℓ -model in H , and hence in G , can be computed in $\mathcal{O}(\ell(|H| + \|H\|))$ time. \square

4. Proof of Theorem 1.2

Let G and H be graphs. An H -partition of G is a proper partition $\{S_v \subseteq V(G) : v \in V(H)\}$ of $V(G)$ indexed by the vertices of H , such that for all distinct $v, w \in V(H)$, we have $vw \in E(H)$ if and only if there is an edge of G between S_v and S_w . Let G_v denote the induced subgraph $G[S_v]$ for each $v \in V(H)$. An H -partition of G is *knitted* if for all distinct $v, w \in V(H)$, we have $vw \in E(H)$ if and only if there is an edge of G between each component of G_v and each component of G_w , as illustrated in Figure 3.

The following lemma, proved shortly, is the heart of the proof of our main result (Theorem 1.2).

LEMMA 4.1. *There is an algorithm that, given $\ell, k \in \mathbb{Z}^+$ and a graph G , outputs a knitted H -partition of G in time $\mathcal{O}(2^{2\ell} \cdot |G| + \|G\|)$, such that either:*

- (a) H contains a K_ℓ -model (which is also output), or
- (b) $|H| \leq 2^{\ell^2+2} \cdot |G| \cdot k^{-1}$, and $|G_x| < 2k$ for all $x \in V(H)$.

Recall the main result of the article.

THEOREM 1.2. *There is an algorithm that, given $\epsilon \in [0, \frac{1}{2}]$, $\ell \in \mathbb{Z}^+$, and a weighted graph (G, w) , outputs either:*

- (a) a K_ℓ -model of G , or
- (b) a $\frac{2}{3}$ -separation of (G, w) of order at most $\ell^{3/2} \cdot 2^{(\ell^2+4)/2} \cdot |G|^{(2-\epsilon)/3}$ in time $\mathcal{O}(\ell \cdot 2^{(3\ell^2+2\ell+6)/2} \cdot |G|^{1+\epsilon} + \ell \cdot \|G\|)$.

PROOF OF THEOREM 1.2 ASSUMING LEMMA 4.1. Apply Lemma 4.1 with $k = \lfloor |G|^{(1-2\epsilon)/3} \rfloor$. We obtain a knitted H -partition of G .

First suppose that case (a) in Lemma 4.1 holds. Thus H contains a K_ℓ -model $\{S_1, S_2, \dots, S_\ell\}$, where each S_i is a connected subgraph of H . Choose a connected component Z_v of G_v for each $v \in V(H)$. For $i \in [\ell]$, let T_i be the induced subgraph $G[\bigcup\{V(Z_v) : v \in V(S_i)\}]$. Since the S_i subgraphs are pairwise disjoint, the T_i subgraphs are pairwise disjoint. Since each S_i is connected in H and each Z_v is connected in G , each T_i subgraph is connected in G . Since the S_i subgraphs are pairwise adjacent, $\{T_1, T_2, \dots, T_\ell\}$ is a K_ℓ -model of G , and case (a) in Theorem 1.2 is satisfied.

Now assume that case (b) in Lemma 4.1 holds. Then

$$|H| \leq 2^{\ell^2+2} \cdot |G| \cdot k^{-1} \leq 2^{\ell^2+2} \cdot |G|^{2(1+\epsilon)/3},$$

and for all $x \in V(H)$,

$$|G_x| < 2k \leq 2|G|^{(1-2\epsilon)/3}.$$

Let $w(v) := w(G_v)$ for all $v \in V(H)$. Apply Theorem 2.3 to (H, w) . The time complexity is

$$\begin{aligned} \mathcal{O}(\ell \cdot 2^\ell \cdot |H|^{3/2} + \ell \cdot \|H\|) &\subseteq \mathcal{O}(\ell \cdot 2^\ell \cdot (2^{\ell^2+2} \cdot |G|^{2(1+\epsilon)/3})^{3/2} + \ell \cdot \|G\|) \\ &\subseteq \mathcal{O}(\ell \cdot 2^{(3\ell^2+2\ell+6)/2} \cdot |G|^{1+\epsilon} + \ell \cdot \|G\|). \end{aligned}$$

We obtain either a K_ℓ -model of H , or a $\frac{2}{3}$ -separation of H with order at most $\ell^{3/2} \cdot |H|^{1/2}$. In the first case, G contains a K_ℓ -model as proved before, and we are done.

Now assume that Theorem 2.3 gives a $\frac{2}{3}$ -separation $\{A, B\}$ of (H, w) with order

$$\begin{aligned} |A \cap B| &\leq \ell^{3/2} \cdot |H|^{1/2} \leq \ell^{3/2} \cdot (2^{\ell^2+2} \cdot |G|^{2(1+\epsilon)/3})^{1/2} \\ &\leq \ell^{3/2} \cdot 2^{(\ell^2+2)/2} \cdot |G|^{(1+\epsilon)/3}. \end{aligned}$$

Let $X := \bigcup\{V(G_v) : v \in A\}$ and $Y := \bigcup\{V(G_v) : v \in B\}$. Then $\{X, Y\}$ is a separation of G . Since $|G_v| < 2|G|^{(1-2\epsilon)/3}$ the order of this separation is

$$\begin{aligned} |X \cap Y| &= \sum_{v \in A \cap B} |G_v| \leq \ell^{3/2} \cdot 2^{(\ell^2+2)/2} \cdot |G|^{(1+\epsilon)/3} \cdot 2|G|^{(1-2\epsilon)/3} \\ &\leq \ell^{3/2} \cdot 2^{(\ell^2+4)/2} \cdot |G|^{(2-\epsilon)/3}. \end{aligned}$$

We have $w(X - Y) = w(A - B) \leq \frac{2}{3}w(H) = \frac{2}{3}w(G)$. Similarly $w(B - A) \leq \frac{2}{3}w(G)$. Therefore $\{X, Y\}$ is a $\frac{2}{3}$ -separation of G . \square

It remains to prove Lemma 4.1.

PROOF OF LEMMA 4.1. Step 1. Initial Partition: Using a linear-time breadth-first search algorithm, compute a maximal set \mathcal{A} of pairwise disjoint subsets of $V(G)$, such that $G[S]$ is connected and $|S| = k$ for each $S \in \mathcal{A}$. Let \mathcal{B} be the set of vertex sets of the connected components of $G - \bigcup\{S : S \in \mathcal{A}\}$. Then $\mathcal{A} \cup \mathcal{B}$ is a partition of $V(G)$.

Step 2. Constuction of H : Let H be the graph such that $\mathcal{A} \cup \mathcal{B}$ is an H -partition of G . Since G_v is connected for each $v \in V(H)$, this H -partition is knitted. Let $A := \{v \in V(H) : V(G_v) \in \mathcal{A}\}$ and $B := \{v \in V(H) : V(G_v) \in \mathcal{B}\}$. A vertex v of

H is big if $|G_v| \geq k$. A vertex v of H is small if $|G_v| < k$. By construction, every vertex in A is big, B is an independent set of H , and every vertex in B is small.

Step 3. Partition of B : Partition $B = C \cup D \cup E$ as follows.

$$\begin{aligned} C &:= \{v \in B : \deg_H(v) \geq 2^{\ell-2}\} \\ D &:= \{v \in B : \ell - 1 \leq \deg_H(v) < 2^{\ell-2}\} \\ E &:= \{v \in B : \deg_H(v) \leq \ell - 2\} \end{aligned}$$

Suppose that $|C| \geq |A|$. Then $H[C \cup A]$ has at least $2^{\ell-2} \cdot |C|$ edges and at most $2|C|$ vertices. By Theorem 2.1, a K_ℓ -model of $H[C \cup A]$ can be computed in $\mathcal{O}(\ell \cdot |G|)$ time, and we are done. Now assume that $|C| < |A|$.

Step 4. Assignment: “Assign” vertices in $D \cup E$ to pairs of distinct vertices in A as follows. Let $\binom{A}{2} := \{\{x, y\} : x, y \in A \text{ and } x \neq y\}$ be the set of pairs of distinct vertices in A . Let Q be the bipartite graph with vertex set $V(Q) := \binom{A}{2} \cup (D \cup E)$, where $\{x, y\} \in \binom{A}{2}$ is adjacent to $v \in D \cup E$ in Q if and only if $x, y \in N_H(v)$. Since each vertex in $D \cup E$ has degree at most $2^{\ell-2}$ in H , each vertex in $D \cup E$ has degree at most $2^{2\ell-4}$ in Q , and Q can be constructed in $\mathcal{O}(2^{2\ell} \cdot |G|)$ time.

Now apply the following greedy algorithm to construct a maximal matching M in Q . (M need not be maximum.) Formally, M is a partial function from $V(Q)$ to $E(Q)$, with M initially undefined everywhere. For each vertex $v \in D \cup E$ in arbitrary order, if v is incident to an edge $\{\{x, y\}, v\} \in E(Q)$, such that no edge in M is incident to $\{x, y\}$, then add (one such edge) $\{\{x, y\}, v\}$ to M . Formally, if $M(\{x, y\})$ is undefined for some edge $e = \{\{x, y\}, v\} \in E(Q)$, then set $M(\{x, y\}) := M(v) := e$. We say that v is assigned to the pair $\{x, y\}$. Since each vertex in $D \cup E$ has degree at most $2^{2\ell-4}$ in Q , this step can be implemented in $\mathcal{O}(2^{2\ell} \cdot |G|)$ time.

Suppose that there is a vertex $v \in D$ that is not assigned; that is, $M(v)$ is undefined. Let $\{x_1, x_2, \dots, x_d\}$ be the neighborhood of v . Then $d \geq \ell - 1$. Thus for all distinct $i, j \in [d]$, there is a distinct vertex $v_{i,j} \in D \cup E$ that is assigned to the pair $\{x_i, x_j\}$, and $v_{i,j}$ is adjacent to both x_i and x_j . In the graph obtained from H by contracting each edge $x_i v_{i,j}$, the subgraph $\{x_1, x_2, \dots, x_d, v\}$ is a clique on $d + 1 \geq \ell$ vertices. Thus H contains a K_ℓ -model, and we are done. This K_ℓ -model can be computed in $\mathcal{O}(2^{2\ell})$ time (since $d < 2^\ell$, and the vertex assigned to a given pair $\{x_i, x_j\}$ can be determined from M in $\mathcal{O}(1)$ time). Hence this step has time complexity $\mathcal{O}(|G| + 2^{2\ell})$. Now assume that every vertex in D is assigned.

Let E^* be the set of assigned vertices in E . Consider the graph obtained from $H[A \cup D \cup E^*]$ by contracting the edge vx for each $v \in D \cup E^*$ assigned to the pair $\{x, y\}$. This graph has $|A|$ vertices and at least $|D| + |E^*|$ edges. Thus if $|D| + |E^*| \geq 2^{\ell-3} \cdot |A|$, then by Theorem 2.1, H contains a K_ℓ -model that can be computed in $\mathcal{O}(\ell \cdot |G|)$ time, and we are done. Now assume that $|D| + |E^*| < 2^{\ell-3} \cdot |A|$.

In total, Step 4 has $\mathcal{O}(2^{2\ell} \cdot |G|)$ time complexity.

Step 5. Handling Unassigned Vertices in E : Partition

$$E - E^* = \bigcup \{P_1, P_2, \dots, P_s\}$$

such that for all $u, v \in E - E^*$, we have $N(u) = N(v)$ if and only if both $u, v \in P_i$ for some $i \in [s]$. By Lemma A.2 in Appendix A, since every vertex in $E - E^*$ has

degree at most $\ell - 2$ in H , this partition can be computed in $\mathcal{O}(\ell \cdot |H|)$ time. For all $i \in [s]$, partition $P_i = \bigcup\{P_{i,1}, P_{i,2}, \dots, P_{i,t_i}\}$ such that

$$k \leq \left| \bigcup\{G_v : v \in P_{i,j}\} \right| < 2k \quad \text{for all } j \in [t_i - 1], \text{ and}$$

$$\left| \bigcup\{G_v : v \in P_{i,t_i}\} \right| < k.$$

This is possible since $|G_v| < k$ for all $v \in P_i$, and can trivially be implemented in $\mathcal{O}(|H|)$ time.

We now determine a new partition of G indexed by a graph H' constructed from H . Collapse each set $P_{i,j}$ of vertices in H into a single vertex $p_{i,j}$ in H' , whose associated subgraph in G is $G_{p_{i,j}} := \bigcup\{G_v : v \in P_{i,j}\}$. The parts A, C, D , and E^* remain unchanged in H' . Since the vertices in $P_{i,j}$ have the same neighborhood, $\{G_v : v \in V(H')\}$ is a knitted partition of G . Let $E_{\text{big}} = \{p_{i,j} : i \in [s], j \in [t_i - 1]\}$ and $E_{\text{small}} = \{p_{i,t_i} : i \in [s]\}$. Then every vertex in E_{big} is big and every vertex in E_{small} is small.

Suppose that $|E_{\text{small}}| \geq 2^{\ell^2} \cdot |A|$. Let X be the graph with vertex set A obtained by adding a clique with vertex set $N_{H'}(v)$ for each vertex $v \in E_{\text{small}}$. Since each such vertex v has degree at most ℓ , the graph X can be constructed in $\mathcal{O}(\ell^2 |H'|)$ time.

We now use this auxiliary graph X to show that, in this case, H' contains a K_ℓ -minor. By construction, X has $|A|$ vertices and at most $\ell^2 \cdot |H|$ edges, and since distinct vertices in E_{small} have distinct neighborhoods, X has at least $|E_{\text{small}}| \geq 2^{\ell^2} \cdot |A|$ cliques. Thus by Lemma 3.2, a K_ℓ -model of X can be computed in time $\mathcal{O}(\ell \cdot (|X| + \|X\|))$, which is $\mathcal{O}(\ell^3 \cdot |H|)$.

For every edge $x_i x_j$ in this K_ℓ -model in X , we have $x_i, x_j \in N(v)$ for some $v \in E_{\text{small}}$. Since v is not assigned, there is a vertex $u \in D \cup E^*$ assigned to $\{x_i, x_j\}$, and u is adjacent to both x_i and x_j . In particular, $M(\{x_i, x_j\}) = \{\{x_i, x_j\}, u\}$ and u can be computed in $\mathcal{O}(1)$ time. Since u is not in the K_ℓ -model, we can include u in the connected subgraph of the K_ℓ -model that contains x_i or x_j , to obtain a K_ℓ -model in $H'[A \cup D \cup E^*]$ (without the edge $x_i x_j$), and we are done. Now assume that $|E_{\text{small}}| < 2^{\ell^2} \cdot |A|$.

In total, Step 5 has time complexity $\mathcal{O}(\ell^2 \cdot |H| + \ell \cdot (|X| + \|X\|)) \leq \mathcal{O}(\ell^3 \cdot |G|)$,

Step 6. Wrapping Up: As illustrated in Figure 4, we have now partitioned $V(H')$ into sets $A \cup E_{\text{big}}$ of big vertices, and sets $C \cup D \cup E^* \cup E_{\text{small}}$ of small vertices. We have proved that $|C| < |A|$, $|D| + |E^*| < 2^{\ell-3} \cdot |A|$, and $|E_{\text{small}}| < 2^{\ell^2} \cdot |A|$. Thus the number of small vertices is less than $(1 + 2^{\ell-3} + 2^{\ell^2}) \cdot |A|$. By definition, the number of big vertices in H' is at most $|G| \cdot k^{-1}$. In particular, $|A| \leq |G| \cdot k^{-1}$. Thus

$$|H'| \leq (1 + 2^{\ell-3} + 2^{\ell^2}) \cdot |A| + |G| \cdot k^{-1} \leq (2 + 2^{\ell-3} + 2^{\ell^2}) \cdot |G| \cdot k^{-1} \leq 2^{\ell^2+2} \cdot |G| \cdot k^{-1}.$$

Moreover, $|H'_v| < 2k$ for every vertex $v \in V(H')$.

The time complexity is $\mathcal{O}(\ell \cdot |G| + \|G\|)$ for Steps 1–3, plus $\mathcal{O}(2^{2\ell} \cdot |G|)$ for Step 4, plus $\mathcal{O}(\ell^3 \cdot |G|)$ for Step 5. Thus the total time complexity is $\mathcal{O}(2^{2\ell} \cdot |G| + \|G\|)$. \square

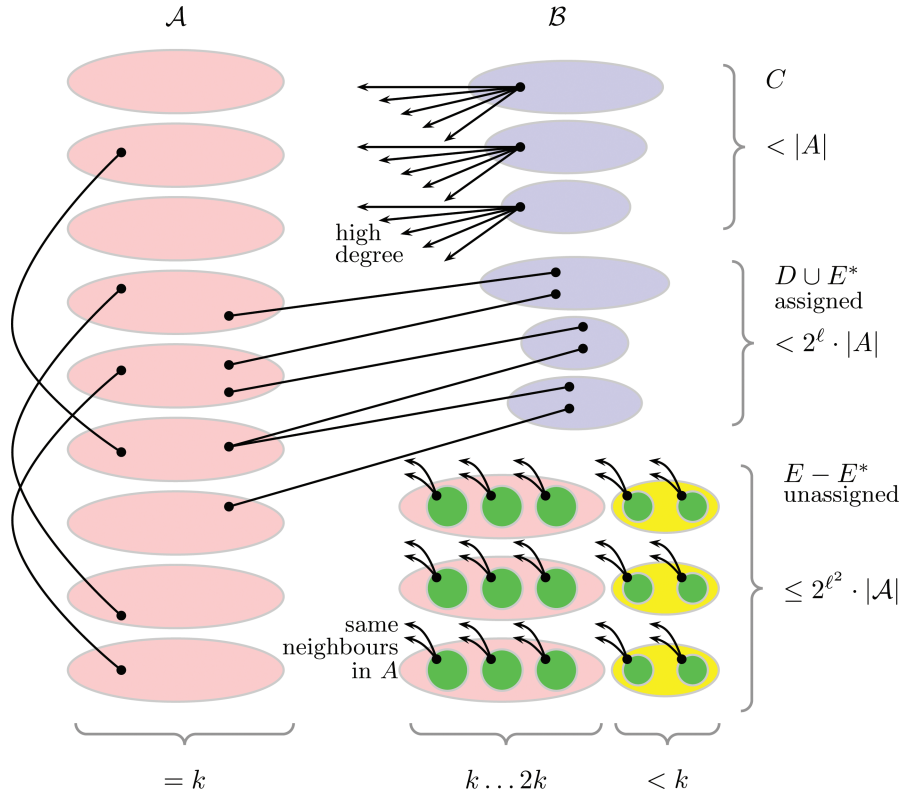


FIG. 4. The partition of $V(G)$ in the proof of Lemma 4.1.

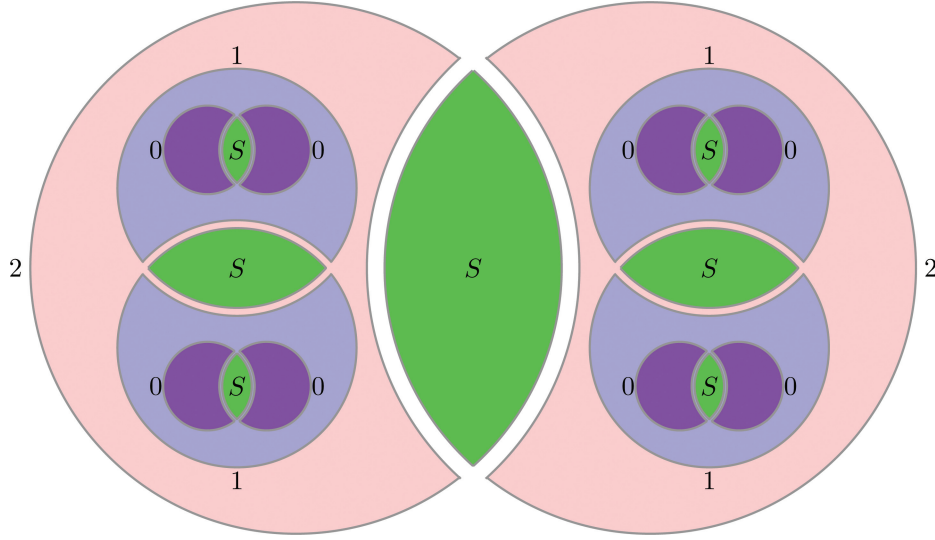
5. Application: Independent Sets

The cardinality of a maximum independent set in a graph G is denoted by $\alpha(G)$. Determining whether $\alpha(G) \geq k$ is a classical \mathcal{NP} -complete problem, and is even hard to approximate in general [Engebretsen and Holmerin 2000; Håstad 1999]. On the other hand, Lipton and Tarjan [1980] showed that separators can be used as the basis for an approximation algorithm for finding independent sets in planar graphs. Using similar ideas, Alon et al. [1990a] outlined an $\mathcal{O}(|G|^{1/2} \cdot \|G\|)$ -time approximation algorithm for finding an independent set in a graph excluding a fixed minor. We improve the time complexity of their algorithm to nearly linear as follows.

THEOREM 5.1. *For fixed ℓ , there is an algorithm that, given a graph G with no K_ℓ -minor, computes an approximation to the maximum independent set of G with relative error $\mathcal{O}((\log \log |G|)^{-1/3})$ in time $\mathcal{O}(|G| \log |G| + \|G\|)$.*

The proof of Theorem 5.1 depends on the following lemma.

LEMMA 5.2. *For fixed ℓ , the following algorithm, given $\epsilon \in [0, 1]$ and a weighted graph (G, w) with no K_ℓ -minor and total weight $w(G) \leq 1$, outputs a set S of $\mathcal{O}(|G|^{2/3} \epsilon^{-1/3})$ vertices of G in time $\mathcal{O}(|G| \log |G| + \|G\|)$, such that every connected component of $G - S$ has weight at most ϵ .*

FIG. 5. Illustration of the computation of S in Lemma 5.2.

```

1: if  $\epsilon \leq |G|^{-1}$  then
2:    $S := V(G)$ 
3: else
4:    $S := \emptyset$ 
5:   while there is a component  $P$  of  $G - S$  with weight exceeding  $\epsilon$  do
6:     let  $\{A, B\}$  be a separation of  $P$  determined by Theorem 1.2 (with  $\epsilon = 0$ )
7:      $S := S \cup (A \cap B)$ 
8:   end while
9: end if
10: output  $S$ 

```

PROOF. If $\epsilon \leq |G|^{-1}$ then $S := V(G)$ satisfies the requirements. Now assume that $\epsilon > |G|^{-1}$. Consider a component P of $G - S$ at some stage of the algorithm. If P is a component of $G - S$ at the termination of the algorithm, then we say P has level 0. Otherwise Theorem 1.2 was applied to P at same stage, to obtain a separation $\{A, B\}$ of P . Thus $w(A - B) \leq \frac{2}{3}w(P)$ and $w(B - A) \leq \frac{2}{3}w(P)$. Each component of $P - (A \cap B)$ is also a component of $G - S$ at some stage of the algorithm. Define the level of P to be 1 plus the maximum level of a component of $P - (A \cap B)$. Observe that two components with the same level are disjoint.

Each level 1 component has weight greater than ϵ , and in general, each level- i component has weight at least $(\frac{2}{3})^{i-1}\epsilon$. Since the total weight of G is at most 1, there are at most $(\frac{2}{3})^{i-1}\epsilon^{-1}$ level- i components. Let k be the maximum level. Then $1 \leq (\frac{2}{3})^{k-1}\epsilon^{-1} \leq (\frac{2}{3})^{k-1}|G|$, which implies that $k \leq 1 + \log_{3/2}|G|$. Since the time complexity of Theorem 1.2 is linear for fixed ℓ , and since two components at the same level are disjoint, the total time complexity is $\mathcal{O}(|G| \log |G| + \|G\|)$.

It remains to prove the upper bound on $|S|$. Let P_1, P_2, \dots, P_t be the components at level i . By Theorem 1.2, the number of vertices added to S by splitting P_1, P_2, \dots, P_t is at most $\mathcal{O}(\sum_{j=1}^t |P_j|^{2/3})$. We have $t \leq (\frac{2}{3})^{i-1}\epsilon^{-1}$ and $\sum_{j=1}^t |P_j| \leq |G|$. For fixed t , the sum $\sum_{j=1}^t |P_j|^{2/3}$, subject to $\sum_{j=1}^t |P_j| \leq |G|$,

is maximized when $|P_j| = |G| \cdot t^{-1}$ for all j . Thus

$$\sum_{j=1}^t |P_j|^{2/3} \leq \sum_{j=1}^t (|G| \cdot t^{-1})^{2/3} = t^{1/3} \cdot |G|^{2/3} \leq \left(\left(\frac{2}{3}\right)^{i-1} \epsilon^{-1}\right)^{1/3} \cdot |G|^{2/3}.$$

Hence

$$|S| \in \mathcal{O}\left(\sum_{i=1}^k \left(\frac{2}{3}\right)^{(i-1)/3} \cdot \epsilon^{-1/3} \cdot |G|^{2/3}\right) \subseteq \mathcal{O}(|G|^{2/3} \epsilon^{-1/3}). \quad \square$$

PROOF OF THEOREM 5.1. Apply Lemma 5.2 with $\epsilon := (\log_2 \log_2 |G|) \cdot |G|^{-1}$, and with each vertex having weight $|G|^{-1}$. We obtain a set S of $\mathcal{O}(|G| \cdot (\log \log |G|)^{-1/3})$ vertices of G such that every component of $G - S$ has weight at most ϵ ; that is, every component of $G - S$ has at most $\log_2 \log_2 |G|$ vertices. In each component of $G - S$, find a maximum independent set by checking every subset of the vertices. Let I be the union of the independent sets obtained. Then I is an independent set of G .

The restriction of a maximum independent set of G to a component of $G - S$ is at most as large as the restriction of I to the same component. Thus

$$\alpha(G) - |I| \leq |S| \in \mathcal{O}(|G| \cdot (\log \log |G|)^{-1/3}).$$

Duchet and Meyniel [1982] proved that $\alpha(G) \geq |G|/2\ell$. Thus the relative error $(\alpha(G) - |I|)/\alpha(G) \in \mathcal{O}((\log \log |G|)^{-1/3})$.

The computation of S takes $\mathcal{O}(|G| \log |G| + \|G\|)$ time by Lemma 5.2.

For each component P of $G - S$ the second step of the algorithm takes $\mathcal{O}(|P| \cdot 2^{|P|})$ time. Thus in total, the second step takes $\mathcal{O}(\sum_P |P| \cdot 2^{|P|})$ time, which is maximized when all components P have the same maximal number of vertices; that is, when $|P| = \log_2 \log_2 |G|$. Hence the second step takes $\mathcal{O}(|G| \cdot 2^{|P|}) = \mathcal{O}(|G| \log |G|)$ time. \square

Appendix

A. More Algorithmic Details

This appendix provides details for some elementary algorithms used in the article.

LEMMA A.1. *The following algorithm, given a graph G that is not d -degenerate (for some $d \in \mathbb{R}^+$), outputs a subgraph H of G in time $\mathcal{O}(|G| + \|G\|)$, such that H has minimum degree greater than d .*

```

1: while there is a vertex  $v$  of degree at most  $d$  in  $G$  do
2:   delete  $v$  from  $G$ 
3: end while
4: output  $G$ 

```

PROOF. The assumption that G is not d -degenerate means that some subgraph of G has minimum degree greater than d . The algorithm finds such a subgraph since a vertex of degree at most d is in no subgraph of G with minimum degree greater than d . Thus upon termination of the algorithm, the remaining subgraph has minimum degree greater than d .

The algorithm can be implemented in $\mathcal{O}(|G| + \|G\|)$ time by maintaining the degree of each vertex in the current graph, and by maintaining a set L of vertices with degree at most d (represented as a boolean function that indicates whether a given vertex is in L in $\mathcal{O}(1)$ time). Clearly L can be initialized in $\mathcal{O}(|G| + \|G\|)$ time. When deleting a vertex v from G , only a neighbor of v needs its degree to be updated, and only a neighbor of v might need to be added to L . Thus when deleting v , these data structures can be maintained in $\mathcal{O}(\deg(v))$ time. Thus the total time complexity is $\mathcal{O}(|G| + \|G\|)$. \square

LEMMA A.2. *There is an algorithm that takes as input a graph G and a set $X \subseteq V(G)$ with $\deg(v) \leq k$ for every vertex $v \in X$, and outputs a partition S_1, \dots, S_k of X such that $v, w \in S_i$ if and only if $N(v) = N(w)$ for all $i \in [k]$. The time complexity is $\mathcal{O}(k \cdot |X|)$.*

PROOF. The following algorithm determines a partial function $f : 2^{V(G)} \rightarrow 2^X$, such that $f(S)$ is defined if and only there is a vertex $v \in X$ with $N_G(v) = S$, and in this case, $f(S) = \{v \in X : N_G(v) = S\}$. The set T is the set of all sets $S \subset V(G)$ for which $f(S)$ is defined.

```

1:  $T := \emptyset$ 
2: for each vertex  $v \in X$  do
3:    $S := N_G(v)$ 
4:   if  $f(S)$  is defined then
5:      $f(S) := f(S) \cup \{v\}$ 
6:   else
7:      $T := T \cup \{S\}$ 
8:      $f(S) := \{v\}$ 
9:   end if
10: end for
11: for  $S \in T$  do
12:   output  $f(S)$ 
13: end for

```

Since $\deg(v) \leq k$ for every vertex $v \in X$, we have $|S| \leq k$, and thus it takes $\mathcal{O}(k)$ time to execute each command inside the loops. The inner steps of each loop are executed $\mathcal{O}(|X|)$ times. Thus the total time complexity is $\mathcal{O}(k \cdot |X|)$. \square

ACKNOWLEDGMENTS. Thanks to the referee whose comments improved the presentation of the article. Thanks to Vida Dujmović for instructive conversations.

REFERENCES

- ALBER, J., FERNAU, H., AND NIEDERMEIER, R. 2003. Graph separators: A parameterized view. *J. Comput. System Sci.* 67, 4, 808–832.
- ALEKSANDROV, L., DJIDJEV, H., GUO, H., AND MAHESHWARI, A. 2006. Partitioning planar graphs with costs and weights. *ACM J. Exp. Algor.* 11.
- ALEKSANDROV, L. G., AND DJIDJEV, H. N. 1989. Improved bounds on the size of separators of toroidal graphs. In *Optimal Algorithms*. Lecture Notes in Computer Science, vol. 401. Springer, 126–138.
- ALEKSANDROV, L. G., AND DJIDJEV, H. N. 1996. Linear algorithms for partitioning embedded graphs of bounded genus. *SIAM J. Discrete Math.* 9, 1, 129–150.
- ALON, N., SEYMOUR, P., AND THOMAS, R. 1994. Planar separators. *SIAM J. Discrete Math.* 7, 2, 184–193.
- ALON, N., SEYMOUR, P. D., AND THOMAS, R. 1990a. A separator theorem for graphs with an excluded minor and its applications. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC'90)*. ACM Press, 293–299.

- ALON, N., SEYMOUR, P. D., AND THOMAS, R. 1990b. A separator theorem for nonplanar graphs. *J. Amer. Math. Soc.* 3, 4, 801–808.
- AMIR, E., KRAUTHGAMER, R., AND RAO, S. 2003. Constant factor approximation of vertex-cuts in planar graphs. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC'03)*. ACM, 90–99.
- ARORA, S., RAO, S., AND VAZIRANI, U. 2004. Expander flows, geometric embeddings and graph partitioning. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC'04)*. ACM, 222–231.
- BODLAENDER, H. L., GILBERT, J. R., HAFSTEINSSON, H., AND KLOKS, T. 1995. Approximating treewidth, pathwidth, frontsize, and shortest elimination tree. *J. Algos.* 18, 2, 238–255.
- CHIBA, N., NISHIZEKI, T., AND SAITO, N. 1981. Applications of the Lipton and Tarjan planar separator theorem. *J. Inf. Process.* 4, 4, 203–207.
- CHUNG, F. R. K. 1991. Improved separators for planar graphs. In *Graph Theory, Combinatorics, and Applications, Vol. 1 (1988)*. Wiley, 265–270.
- DEMAINE, E. D., AND HAJIAGHAYI, M. 2005. Graphs excluding a fixed minor have grids as large as treewidth, with combinatorial and algorithmic applications through bidimensionality. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*. ACM, 682–689.
- DEMAINE, E. D., AND HAJIAGHAYI, M. 2008a. The bidimensionality theory and its algorithmic applications. *The Comput. J.* 51, 3, 292–302.
- DEMAINE, E. D., AND HAJIAGHAYI, M. 2008b. Linearity of grid minors in treewidth with applications through bidimensionality. *Combinatorica* 28, 1, 19–36.
- DIKS, K., DJIDJEV, H. N., SÝKORA, O., AND VRTO, I. 1993. Edge separators of planar and outerplanar graphs with applications. *J. Algor.* 14, 2, 258–279.
- DJIDJEV, H. N. 1981. A separator theorem. *C. R. Acad. Bulgare Sci.* 34, 5, 643–645.
- DJIDJEV, H. N. 1982. On the problem of partitioning planar graphs. *SIAM J. Algebraic Discrete Methods* 3, 2, 229–240.
- DJIDJEV, H. N. 1985a. A linear algorithm for partitioning graphs of fixed genus. *Serdica* 11, 4, 369–387.
- DJIDJEV, H. N. 1985b. A separator theorem for graphs of fixed genus. *Serdica* 11, 4, 319–329.
- DJIDJEV, H. N. 1987. On the constants of separator theorems. *C. R. Acad. Bulgare Sci.* 40, 10, 31–34.
- DJIDJEV, H. N. 2000. Partitioning planar graphs with vertex costs: Algorithms and applications. *Algorithmica* 28, 1, 51–75.
- DJIDJEV, H. N., AND GILBERT, J. R. 1999. Separators in graphs with negative and multiple vertex weights. *Algorithmica* 23, 1, 57–71.
- DJIDJEV, H. N., AND VENKATESAN, S. M. 1997. Reduced constants for simple cycle graph separation. *Acta Inf.* 34, 3, 231–243.
- DUCHET, P., AND MEYNIEL, H. 1982. On Hadwiger's number and the stability number. *Ann. Discrete Math.* 13, 71–73.
- DUJMOVIĆ, V., AND WOOD, D. R. 2004. Three-Dimensional grid drawings with sub-quadratic volume. In *Towards a Theory of Geometric Graphs*, J. Pach, Ed. Contemporary Mathematics, vol. 342. American Mathematics Society, 55–66.
- ENGBRETSSEN, L., AND HOLMERIN, J. 2000. Clique is hard to approximate within $n^{1-\alpha(1)}$. In *Proceedings of the 27th International Colloquium on Automata, Languages and Programming (ICALP'00)*. Lecture Notes in Comput. Science., vol. 1853. Springer, 2–12.
- EVEN, G., NAOR, J., RAO, S., AND SCHIEBER, B. 1999. Fast approximate graph partitioning algorithms. *SIAM J. Comput.* 28, 6, 2187–2214.
- EVEN, G., NAOR, J., RAO, S., AND SCHIEBER, B. 2000. Divide-and-Conquer approximation algorithms via spreading metrics. *J. ACM* 47, 4, 585–616.
- FEIGE, U., AND MAHDIAN, M. 2006. Finding small balanced separators. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC'06)*. ACM, 375–384.
- GARG, N., SARAN, H., AND VAZIRANI, V. V. 1999. Finding separator cuts in planar graphs within twice the optimal. *SIAM J. Comput.* 29, 1, 159–179.
- GILBERT, J. R., HUTCHINSON, J. P., AND TARJAN, R. E. 1984. A separator theorem for graphs of bounded genus. *J. Algor.* 5, 3, 391–407.
- GROHE, M. 2003. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica* 23, 4, 613–632.
- HÅSTAD, J. 1999. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Math.* 182, 1, 105–142.
- KOSTOCHKA, A. V. 1982. The minimum Hadwiger number for graphs with a given mean degree of vertices. *Metody Diskret. Analiz.* 38, 37–58.

- KOSTOCHKA, A. V. 1984. Lower bound of the Hadwiger number of graphs by their average degree. *Combinatorica* 4, 4, 307–316.
- LEISERSON, C. E. 1980. Area-Efficient graph layouts (for VLSI). In *Proceedings of the 21st Annual Symposium on Foundations of Computer Science (FOCS'80)*. IEEE, 270–281.
- LIPTON, R. J., ROSE, D. J., AND TARJAN, R. E. 1979. Generalized nested dissection. *SIAM J. Numer. Anal.* 16, 2, 346–358.
- LIPTON, R. J., AND TARJAN, R. E. 1979. A separator theorem for planar graphs. *SIAM J. Appl. Math.* 36, 2, 177–189.
- LIPTON, R. J., AND TARJAN, R. E. 1980. Applications of a planar separator theorem. *SIAM J. Comput.* 9, 3, 615–627.
- MADER, W. 1967. Homomorphieeigenschaften und mittlere kantendichte von graphen. *Math. Ann.* 174, 265–268.
- MILLER, G. L. 1986. Finding small simple cycle separators for 2-connected planar graphs. *J. Comput. System Sci.* 32, 3, 265–279.
- MILLER, G. L., TENG, S.-H., THURSTON, W., AND VAVASIS, S. A. 1997. Separators for sphere-packings and nearest neighbor graphs. *J. ACM* 44, 1, 1–29.
- NORINE, S., SEYMOUR, P., THOMAS, R., AND WOLLAN, P. 2006. Proper minor-closed families are small. *J. Combin. Theory Ser. B* 96, 5, 754–757.
- PLAISTED, D. A. 1990. A heuristic algorithm for small separators in arbitrary graphs. *SIAM J. Comput.* 19, 2, 267–280.
- PLOTKIN, S., RAO, S., AND SMITH, W. D. 1994. Shallow excluded minors and improved graph decompositions. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'94)*. ACM, 462–470.
- ROBERTSON, N., AND SEYMOUR, P. D. 1995. Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory Ser. B* 63, 1, 65–110.
- ROSENBERG, A. L., AND HEATH, L. S. 2001. *Graph Separators, with Applications*. Frontiers of Computer Science. Kluwer.
- SMITH, W. D., AND WORMALD, N. C. 1998. Geometric separator theorems and applications. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science (FOCS'98)*. IEEE, 232–243.
- ŠYKORA, O., AND VŘÍTO, I. 1993. Edge separators for graphs of bounded genus with applications. *Theoret. Comput. Sci.* 112, 2, 419–429.
- TAZARI, S., AND MÜLLER-HANNEMANN, M. 2009. Shortest paths in linear time on minor-closed graph classes, with an application to Steiner tree approximation. *Discrete Appl. Math.* 157, 4, 673–684.
- THOMASON, A. 1984. An extremal function for contractions of graphs. *Math. Proceedings of the Cambridge Philos. Soc.* 95, 2, 261–265.
- THOMASON, A. 2001. The extremal function for complete minors. *J. Combin. Theory Ser. B* 81, 2, 318–338.
- VENKATESAN, S. M. 1987. Improved constants for some separator theorems. *J. Algor.* 8, 4, 572–578.
- WOOD, D. R. 2007. On the maximum number of cliques in a graph. *Graphs Combin.* 23, 3, 337–352.
- YUSTER, R. 2008. Single source shortest paths in H -minor free graphs. arXiv:0809.2970.
- YUSTER, R., AND ZWICK, U. 2007. Maximum matching in graphs with an excluded minor. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'07)*. SIAM, 108–117.

RECEIVED APRIL 2005; REVISED APRIL 2009; ACCEPTED JUNE 2009