A SEPARATOR THEOREM FOR
GRAPHS OF BOUNDED GENUS

John R. Gilbert^{*}
Joan P. Hutchinson
Robert Endre Tarjan

TR 82-506
July 1982

Department of Computer Science
Upson Hall
Cornell University
Ithaca, New York 14853

^{*}The work of this author was supported in part by National Science Foundation grant MCS-82-02948.

A Separator Theorem for Graphs of Bounded Genus

John R. Gilbert^{*}
Computer Science Department
Cornell University
Ithaca, New York

Joan P. Hutchinson
Mathematics Department
Smith College
Northampton, Massachusetts

Robert Endre Tarjan
Bell Laboratories
Murray Hill, New Jersey

July 1982

Abstract

Many divide-and-conquer algorithms on graphs are based on finding a small set of vertices or edges whose removal divides the graph roughly in half. Most graphs do not have the necessary small separators, but some useful classes do. One such class is planar graphs: If we can draw an n -vertex graph on the plane, then we can bisect it by removing $O(\sqrt{n})$ vertices [Lipt79b].

The main result of this paper is that if we can draw a graph on a surface of genus g , then we can bisect it by removing $O(\sqrt{gn})$ vertices. This bound is best possible to within a constant factor. We give an algorithm for finding the separator that takes time linear in the number of edges in the graph, given an embedding of the graph in its genus surface. We discuss some extensions and applications of these results.

^{*} The work of this author was supported in part by National Science Foundation grant MCS-82-02948.

I. Introduction

Many divide-and-conquer algorithms on graphs are based on finding a small set of vertices or edges whose removal divides the graph roughly in half. Examples include layout of circuits in a model of VLSI [Leis80], efficient sparse Gaussian elimination [Lipt79a, Gilb80], and construction of Voronoi diagrams to solve various geometric problems [Lipt80].

A class of graphs is said to satisfy an $f(n)$ -separator theorem if there is a constant $\alpha < 1$ such that every n -vertex graph in the class has a set of $O(f(n))$ vertices^{*} whose removal leaves no connected component with more than αn vertices. Sparse graphs in general do not satisfy a nontrivial separator theorem [Lipt79b], but several useful classes of such graphs do. Separator theorems are known for trees [Lew65], outerplanar graphs [Leis80], grid graphs [Lipt79b], hypercubes [Gilb80], several interconnection graphs for parallel computation [Hoey80, Leis80, Leig81], planar graphs [Lipt79b], and graphs of bounded genus [Lipt80]. This paper concerns the last class.

Lipton and Tarjan proved a separator theorem for planar graphs.

Theorem 1. [Lipt79b] A planar graph with n vertices has a set of at most $2\sqrt{2n}$ vertices whose removal leaves no component with more than $2n/3$ vertices. $\square$

* We use the following asymptotic notation: $f(n) = O(g(n))$ means there are constants c and m such that $n > m$ implies $|f(n)| \leq |cg(n)|$; $f(n) = o(g(n))$ means for all $c > 0$ there exists an m such that $n > m$ implies $|f(n)| \leq |cg(n)|$; $f(n) = \Omega(g(n))$ means $g(n) = O(f(n))$; and $f(n) = \Theta(g(n))$ means $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$.

This bound is tight to within a constant factor. Djidjev [Djid82] improved the constant $2\sqrt{2}$ to $\sqrt{6}$ and proved that the constant must be at least $\sqrt{4\pi\sqrt{3}}/3$; the best possible constant is not known. Lipton and Tarjan also gave an $O(n)$ -time algorithm to find the separator.

The genus of a graph is the minimum number of handles that must be added to a sphere so that the graph can be embedded in the resulting surface with no crossing edges. Albertson and Hutchinson proved that the genus of a graph can be reduced by removing a small set of vertices.

Theorem 2. [Albe78] A graph of genus $g > 0$ with n vertices has a set of at most $\sqrt{2n}$ vertices whose removal leaves a graph of genus at most $g-1$. $\square$

This immediately gives a separator theorem for graphs of genus g .

Theorem 3. [Lipt80] A graph of genus g with n vertices has a set of at most $(g+2)\sqrt{2n}$ vertices whose removal leaves no component with more than $2n/3$ vertices. $\square$

The main result of this paper is the following improvement of Theorem 3.

Theorem 4. A graph of genus g with n vertices has a set of at most $6\sqrt{gn} + 2\sqrt{2n} + 1$ vertices whose removal leaves no component with more than $2n/3$ vertices. $\square$

The next section presents a proof of Theorem 4. Section III shows that this result is tight to within a constant factor. Section IV presents an algorithm that finds a $\sqrt{gn}$ -separator in time $O(n+g)$ (which is of the order of the number of edges in the graph). The input to the algorithm is an embedding of the graph in its genus surface. Section V describes applications, extensions, and open problems.

II. A $\sqrt{gn}$ -separator theorem

Before proving Theorem 4, we shall state some definitions and results from topological graph theory. We refer the reader to Massey [Mas67] and White [Whit73] for proofs and more details.

Let a graph G be embedded in an orientable surface S of genus g . The faces of G (more properly, of the embedding of G) are the connected components of $S-G$. A face that is simply connected (that is, homeomorphic to an open disc) is called a 2-cell. If every face of G is a 2-cell, the embedding is a 2-cell embedding. If G can be embedded in a surface S of genus g but not in a surface of genus $g-1$, we say that G has genus g and that S is its genus surface.

Figures 1 and 2 show examples of embeddings in a torus. In each figure the top and bottom borders of the rectangle are identified, as are the left and right borders. Figure 1 shows two non-2-cell embeddings of a triangle in a torus. In each, one face is not simply connected. Figure 2 shows 2-cell toroidal embeddings of the complete graphs on four and five vertices.

[Figure 1]

[Figure 2]

Theorem 5. Every embedding of a connected graph in its genus surface is a 2-cell embedding. $\square$

The converse of this theorem is false. For example, the complete graph on four vertices, which is planar, has a 2-cell embedding in a torus.

Theorem 6. (Euler's formula) Let G be a connected graph with n vertices and m edges. If G has a 2-cell embedding with f faces in a surface of genus g , then $n-m+f = 2-2g$. $\square$

A contraction of a graph is obtained by repeating the following step as many times as desired: shrink an edge $\{v,w\}$ into a single vertex adjacent to all vertices originally adjacent to v or w .

Theorem 7. If G has genus g , then any contraction of G has genus at most g . $\square$

This is all the machinery we need to prove Theorem 4. We shall trade off the size of a separating set, which can be as big as $\sqrt{gn}$, for the size of a genus-reducing set, which is limited to $\sqrt{n/g}$. We first reduce the problem to a subgraph with a spanning tree of depth $\sqrt{n/g}$. Then we find a set of $O(\sqrt{gn})$ vertices of that tree whose removal makes the subgraph planar. This set is the union of $O(g)$ paths in the tree. Finally we separate the planar subgraph by removing $O(\sqrt{n})$ more vertices.

Proof of Theorem 4. Let G be a connected graph of genus g with n vertices. We shall find a set C of at most $6\sqrt{gn} + 2\sqrt{2n} + 1$ vertices whose removal leaves no connected component of G with more than $2n/3$ vertices. Since the planar separator theorem gives C if $g=0$, we assume $g>0$.

1. Reduction to a graph with a short spanning tree

Choose a vertex x of G . Partition the vertices into levels $L_0, L_1, \dots, L_k$ according to their distance from x ; let L_{k+1} be an empty highest level. Let L be a level such that at most $n/2$ vertices are on levels lower than L and at most $n/2$ vertices are on levels higher than L . If $|L| \leq \sqrt{gn}$ then the vertices on level L are the required separator. Otherwise, let M be the lowest level above L such that $|M| \leq \sqrt{gn}$, and let N be the highest level below L such that $|N| \leq \sqrt{gn}$. Removal of M and N divides G into three pieces (some of which may be empty): at most $n/2$ vertices on levels above M , at most $n/2$ vertices on levels below N , and the vertices on the levels between N and M . There are fewer than $\sqrt{n/g}$ levels between N and M , since each such level contains more than $\sqrt{gn}$ vertices.

Discard the vertices on levels M and above, and contract the vertices on levels N and below into a single vertex r . Call the resulting graph H . The genus h of H is at most g , and H has a spanning tree rooted at r with depth less than $\sqrt{n/g}$. Each edge of H is either a tree edge or a non-tree edge.

Consider an embedding of H in its genus surface S' . Suppose H has v vertices, e edges, and f faces. Since the embedding is 2-cell, $v - e + f = 2 - 2h$.

2. Making H planar

The next step is to find a set of $O(\sqrt{gn})$ vertices whose removal makes H planar. We will do this backwards, by throwing away all but $O(\sqrt{gn})$ of the vertices of H and then arguing that the vertices we threw away induce a planar subgraph of H . Figure 3 is an example of a graph H embedded in a torus. In Figure 3(a), tree edges are solid lines and non-tree

edges are dotted lines.

[Figure 3(a)]

First we delete non-tree edges one at a time until H has only one face. We choose which edges to delete so that after each deletion the faces of H all remain 2-cells, as follows. Suppose that H has more than one face. Let F be a face. We claim that there is a non-tree edge $\{a,b\}$ that separates F from some other face.

To see this, color F red and the remaining faces blue. The boundary between the red and blue regions of S' is made up of edges of H . Each vertex that is an endpoint of a boundary edge is the endpoint of at least two boundary edges. Choose a boundary edge between red and blue. Construct a path in H by following boundary edges. This path will eventually repeat a vertex. This determines a simple cycle in H each of whose edges has F on one side and a different face on the other side. (The cycle may or may not be the entire boundary of F , since the boundary of F may not be a simple cycle and may include edges that have F on both sides.) The simple cycle must contain a non-tree edge $\{a,b\}$. As desired, $\{a,b\}$ separates F from some other face E .

If we delete $\{a,b\}$ from the embedding, F and E merge into a single face. The resulting face is a 2-cell. Thus Euler's formula still holds (indeed, we have just decreased f and e by one). We continue to delete non-tree edges until the graph has only one face. Figure 3(b) shows the sample graph at this point.

[Figure 3(b)]

Now we repeatedly delete vertices of degree one and their incident edges (which are tree edges). Each such deletion leaves the single face a 2-cell, so it preserves Euler's formula (by decreasing v and e by one). Finally we are left with a connected subgraph H' of H that has one face and no vertex of degree one. See Figure 3(c). If H' has v' vertices and e' edges, we know that $v' - e' + 1 = 2 - 2h$, whence $e' = v' - 1 + 2h$.

[Figure 3(c)]

Now H' is still spanned by a subtree of the original spanning tree of H . Thus H' has $v' - 1$ tree edges and $2h$ non-tree edges. Every tree edge of H' is on the tree path from an endpoint of some non-tree edge to the root r of the spanning tree of H . Since each such tree path has fewer than $\sqrt{n/g}$ edges, H' has fewer than $4h \sqrt{n/g}$ tree edges. This is at most $4 \sqrt{gn}$, so H' has fewer than $4 \sqrt{gn}$ tree edges and hence at most $4 \sqrt{gn} + 1$ vertices.

If we remove H' from H , the rest of H is embedded in the 2-cell that is the single face of H' ; in particular, it is planar.

3. Separating the planar graph

Using Theorem 1, we separate the planar graph $H - H'$ into components, each with at most $2n/3$ vertices, by removing another set P of at most $2 \sqrt{2n}$ vertices. Therefore the original graph G is divided into components, each

with at most $2n/3$ vertices, by removing the vertices in $C = \text{NUMUH}'UP$. There are at most $6 \sqrt{gn} + 2 \sqrt{2n} + 1$ vertices in C . This completes the proof of Theorem 4. $\square$

III. A lower bound

In this section we prove that the bound in Theorem 4 is tight to within a constant factor whenever $g \leq n$. If $g > n$, the trivial n -separator theorem that all graphs satisfy is tighter than the $\sqrt{gn}$ -separator theorem above.

We begin with the case $g = \bar{O}(n)$. We need a graph with n vertices and genus $\bar{O}(n)$ that satisfies no $o(n)$ -separator theorem. This follows from the existence of a class of graphs with $O(n)$ edges and no nontrivial separator theorem. Lipton and Tarjan [Lipt79b] proved nonconstructively that such a class exists. Gilbert exhibited such a class, based on the "expanders" that Gabber and Galil [Gabb79] used in the first explicit construction of sparse superconcentrators.

Theorem 8. [Gilb80] Suppose a constant $\alpha < 1$ and a function $p(n) = o(n)$ are given. For infinitely many n there is an n -vertex graph that is regular of degree three that cannot be divided into pieces as small as αn by removing $p(n)$ vertices. (In fact, for all sufficiently large k , there is such a graph with $3k^2 \leq n \leq 8k^2$.) $\square$

The graphs in Theorem 8 have genus at most $3n/2$, since they can be

embedded by placing all the vertices on a sphere and then adding a handle for each edge. Thus they satisfy no separator theorem better than $\sqrt{gn} = \Theta(n)$.

This proves that Theorem 4 is tight at least for $g = \Theta(n)$. We can show that it is tight for arbitrary $g < n$ by constructing a graph with $\Theta(n)$ vertices and genus $\Theta(g)$ that has no better separator than $\sqrt{gn}$. We start with an expander with $\Theta(g)$ vertices, embedded in a surface of genus $\Theta(g)$. Then we replace each vertex by a pyramid of $k(k+1)/2$ vertices (where k will be specified later), and replace each edge by $2k-1$ edges, as shown in Figure 4. The resulting graph G has $\Theta(k^2 g)$ vertices. Its genus is $\Theta(g)$ since it is embedded in the same surface as the original graph. As we shall see below, separating G requires removing $\Omega(k)$ vertices from each of $\Omega(g)$ pyramids; thus G has no $o(kg)$ -separator theorem.

[Figure 4]

Theorem 9. For any $n > 0$ and $g < n$, there is a graph G with $\Theta(n)$ vertices and genus $\Theta(g)$ that satisfies no $o(\sqrt{gn})$ -separator theorem.

Proof. Let G be as above, taking $k = \sqrt{n/g}$. $\square$

It remains only to show that G has no separator with fewer than kg vertices. The proof is messy but not difficult; it is nearly the same as the proof that a square grid graph has no separator with fewer than $\sqrt{n}$ vertices.

Lemma 1. Let $H = H_k$ be either the $k(k+1)/2$ -vertex pyramid or the $k(k+1)$ -vertex graph consisting of two such pyramids joined by $2k-1$ edges as shown in

Figure 5(b). Let $\alpha < 1$ be given. Then there exists $\beta > 0$ such that any set C whose removal from H leaves no set with more than αn vertices has size at least βk .

Proof. First let H be the $k(k+1)/2$ -vertex pyramid. Define a "file" as the set of vertices in one row and column that intersect on the main diagonal, for example the white vertices in Figure 5(a). There are k files, each file has k vertices, each pair of files intersects in one vertex, and each vertex is in one or two files.

[Figure 5]

Let C be a set of vertices whose removal leaves no component with more than αn vertices, and let A be a component of $H-C$. We claim that either A contains a vertex in every file or $H-A$ contains a vertex in every file. For if A does not, then some file has only vertices of $H-A$; and that includes a vertex from every file.

We consider only the case where A has a vertex in every file; the other case is similar. First we bound the number of files containing only vertices of A . Any p files contain $pk - p(p-1)/2$ vertices, so if A contains p full files then $|A| \leq \alpha n$ implies

$$pk - p(p-1)/2 \leq \alpha k(k+1)/2.$$

We can rearrange this to give

$$(1-\alpha)(k+O(1))^2 \leq (k-p)^2,$$

whence

$$p \leq k(1 - \sqrt{1 - \alpha}) + O(1).$$

Since there are only p files with no vertices of $H-A$, there are $\Omega(k)$ files that contain vertices of $H-A$. Any file that has both a vertex of A and a vertex of $H-A$ must have a vertex of C , so there are $\Omega(k)$ vertices in C .

The proof when H is two joined pyramids is similar. We ignore the first row, and we define a file to be a row and column that intersect on the main diagonal of the remaining square grid. The rest of the proof follows that above, and we omit the details. $\square$

Lemma 2. Fix $\alpha < 1$, and for any positive g and k let $G = G_{g,k}$ be the graph constructed above. Let $n = \Theta(gk^2)$ be the number of vertices in G . Then any set C of vertices of G whose removal leaves no component with more than αn vertices has size $\Omega(gk)$.

Proof. Let C be a set of vertices that divides G into components $A_1, \dots, A_r$, all no larger than αn . Let G' be the expander that results from contracting each pyramid of G to a single vertex. Let $n' = \Theta(g)$ be the number of vertices in G' . We shall use C to find a separator in G' and use the size of that separator to bound the size of C .

By Lemma 1, there is an $\epsilon > 0$ such that (a) removing any set of ϵk vertices from a pyramid leaves a component with at least $\sigma k(k+1)/2$ vertices, where $\sigma = (\alpha+1)/2$, and (b) removing any set of $2\epsilon k$ vertices from a double pyramid (Figure 5(b)) leaves a component with at least $(\sigma+1)k(k+1)/2$ vertices. This implies that if a pyramid contains at most ϵk vertices of C then the majority

of its vertices are in some single component of $G-C$. Let C' be the set of pyramids with more than ϵk vertices of C . Put each remaining pyramid in A'_i if most of its vertices are in A_i . Now C' and $A'_1, \dots, A'_r$ partition the vertices of G' .

First we claim that each A'_i is small. Since A_i has at most αn vertices and each pyramid in A'_i contains at least $\sigma k(k+1)/2 = \sigma n/n'$ of them, there are at most $(\alpha/\sigma)n'$ pyramids in A'_i . Note that α/σ is less than 1.

Next we claim that C' separates A'_i from A'_j for all $i \neq j$. Suppose there is a pyramid P in A'_i adjacent to a pyramid Q in A'_j . By definition of A'_i , neither P nor Q contains ϵk vertices of C . Thus the double pyramid PQ has fewer than $2\epsilon k$ vertices of C , so it has at least $(\sigma+1)k(k+1)/2$ vertices in some component of $G-C$. But then both P and Q have at least $\sigma k(k+1)/2$ vertices in that component, so they are both in the same set A'_i .

Now G' is an n' -vertex expander with a separator C' whose removal leaves no component with more than $\sigma n'$ vertices. Therefore C' has $\Omega(n') = \Omega(g)$ vertices. Each pyramid corresponding to a vertex in C' contains at least ϵk vertices of C , so C has $\Omega(gk)$ vertices, as desired. $\square$

IV. An $O(n+g)$ algorithm

Assume that the connected n -vertex graph G of genus g is given along with an embedding in its genus surface S . This section gives an $O(n+g)$ -time

algorithm to find the separator in Theorem 4. If G has m edges, we know that $n+2g-1 \leq m \leq 3n+6g-6$; thus $m = \Theta(n+g)$.

Any 2-cell embedding can be specified by giving, for each vertex v , the clockwise cyclic ordering of the edges incident with v . (Indeed, every way of ordering the edges corresponds to a 2-cell embedding.) We shall assume that G is represented by a cyclic doubly-linked list, for each vertex, of its incident edges in the order they occur in the embedding. Each edge $\{v, w\}$ occurs twice; we shall consider the occurrence in v 's incidence list to be directed from v to w and the occurrence in w 's list from w to v . Each occurrence of the edge has a pointer to the other. Each edge will be marked (in Step 1 below) as a tree edge or a non-tree edge. Also, each edge has a pointer to the face on its right in the embedding. We shall describe the representation of a face in Step 3, where it is constructed.

The algorithm follows the proof of Theorem 4. We assert that each of the following steps can be carried out in $O(m)$ time.

1. Breadth-first search

Select a root r and carry out a breadth-first search to determine a spanning tree T of G , the levels L_i for $0 \leq i \leq k$, and the separating levels M and N . Let $C = M \cup N$.

2. Delete and contract

Delete vertices on levels N and above, and contract the vertices on levels M and below to the vertex r , forming H . Simultaneously transform the tree T to a spanning tree T' of H .

Specifically, to obtain the embedding of H , we modify the adjacency lists as follows. Let T^* be the subtree of T from r up through the level M ; let M^* be the first level above M . As in the planar separator algorithm [Lipt79b], scan the edges joining a vertex of T^* with a vertex of M^* by traversing clockwise around the tree T^* . Let r 's adjacency list be the vertices of M^* listed in sequential order as met in this traversal, but omitting repeats. For each v in M^* , modify v 's adjacency list by deleting every instance of a vertex in T^* , but replace the first such instance by r . This new representation provides an embedding of H on some surface S' of genus h' . S' may not be the genus surface of H (which has genus h), but $h \leq h' \leq g$. Also, the embedding is 2-cell, which is all that is needed for the rest of the construction in Theorem 4.

3. Identify faces

Each directed edge is considered incident with the face to its right. Select a directed edge e whose incident face is unknown, and use the adjacency lists to trace the boundary of e 's incident face. Repeat this until the incident face of every directed edge is known. Each face is represented by a linked list of its incident edges, and each edge is given a pointer to its incident face. The representation of a face also includes a one-bit mark that will be used in Step 4.

4. Merge faces

Choose one face F to expand by deleting non-tree edges. Each face will be marked "merged" or "unmerged"; initially mark F merged and all other faces unmerged. Consider an edge e in the list representing F 's boundary. If e is

a tree edge, ignore it. Otherwise determine whether the face E incident on the reverse of e is merged or unmerged. If E is merged, ignore e . If E is unmerged, delete e and its reverse from H , append E 's edge list to F 's, and mark E merged. Repeat this for each edge in F 's edge list. This takes $O(m)$ time since each directed edge is considered once.

5. Delete vertices of degree one

Let W be an empty set of vertices. Perform a depth-first search of the spanning tree T' of H . If a vertex has degree one in H after its subtree has been searched, delete it from H and add it to W . After this step, all that remains of H is the planarizing set H' .

6. Add planarizing set to separator

Add the remaining vertices of H to C .

7. Separate planar graph

Apply the planar separator algorithm [Lipt79b] to the graph induced by the vertices in W . Add the resulting separator P to C . (The planar separator algorithm need not find a planar embedding for W ; the embedding is given by the adjacency lists for W 's vertices in H , ignoring vertices not in W .)

V. Remarks

We conclude by mentioning a few extensions, applications, and conjectures. As in the planar case, we can extend Theorem 4 by allowing weights on the vertices and further restricting the sizes of the components remaining after the separator is removed. The following theorems are easily proved by the methods of Lipton and Tarjan [Lipt79b, Lipt80].

Theorem 10. Let G be an n -vertex graph of genus g whose vertices have non-negative weights. The vertices of G can be partitioned into three sets A , B , and C such that no edge joins a vertex in A with a vertex in B , C contains $O(\sqrt{gn})$ vertices, and neither A nor B contains more than half the total weight. The set C can be found in $O(n+g)$ time. $\square$

Theorem 11. Let G be an n -vertex graph of genus g whose vertices have non-negative weights summing to no more than one. Let $0 < \epsilon \leq 1$. Then there is a set C of $O(\sqrt{gn/\epsilon})$ vertices whose removal leaves no connected component with total weight exceeding ϵ . The set C can be found in $O((n+g) \log n)$ time. $\square$

Theorem 2 implies that an n -vertex graph of genus g has a set of $O(g \sqrt{n})$ vertices whose removal leaves a planar graph. We can improve that result as follows. When Theorem 4 is applied to separate a graph, some of the fragments produced may be nonplanar. If we apply Theorem 4 recursively to the nonplanar fragments we get a planar (and probably disconnected) graph after removing $O(\sqrt{gn} \log n)$ vertices. We can improve this further by

selecting our cuts to split the genus rather than the number of vertices in half, giving the following result.

Theorem 12. An n -vertex graph of genus g has a set of $O(\sqrt{gn} \log g)$ vertices whose removal leaves a (possibly disconnected) planar graph.

Remark. We conjecture that this bound can be improved to $O(\sqrt{gn})$.

Proof. Start by partitioning the vertices into levels $L_0, \dots, L_k$ as in the proof of Theorem 4. Removing a level cuts the graph into subgraphs whose genres add to at most g (since the genus of a graph is the sum of the genres of its connected components [Batt62]). Let L_i be the highest level such that the subgraph induced by vertices on levels below L_i has genus at most $g/2$. The subgraph induced by vertices on levels L_{i+2} through L_k has genus less than $g/2$, since the subgraph induced by levels L_0 through L_i has genus more than $g/2$. Thus removing any level at or below L_i and any level at or above L_{i+1} leaves a top piece of genus less than $g/2$, a bottom piece of genus at most $g/2$, and a middle piece.

As in the proof of Theorem 4, let M be the lowest level at or above L_{i+1} with at most $\sqrt{gn}$ vertices, and let N be the highest level at or below L_i with at most $\sqrt{gn}$ vertices. Deleting M and N gives two pieces A and B of genus at most $g/2$ each, and a middle piece with a spanning tree of height less than $\sqrt{n/g}$.

Make the middle piece planar by deleting $4\sqrt{gn}$ vertices, as in the proof of Theorem 4. Apply this whole procedure recursively to make each of A and B planar.

It remains only to count the vertices we have removed. Let $s(n, g)$ be the

maximum number of vertices removed by this procedure from any n -vertex graph of genus g . Then

$$s(n,0) = s(1,g) = 0$$

$$s(n,g) \leq 6\sqrt{gn} + \max \{s(p,g/2) + s(q,g/2) : p+q \leq n\} \text{ for } n > 1, g > 0.$$

The solution to this recurrence is $s(n,g) \leq 6\sqrt{gn} \log_2 g + O(\sqrt{gn})$. $\square$

Most of the applications [Lipt79a, Lipt80] of the planar separator theorem carry over to graphs of genus g in the obvious way. For example, we can find an approximately maximum independent set I in a graph of genus g by applying Theorem 11.

1. Separate

Apply Theorem 11 with $\epsilon = k(n)/n$ (where $k(n)$ is a function to be specified later) to obtain C of size $O(n \sqrt{g/k(n)})$ and components with at most $k(n)$ vertices.

2. Find independent sets

In each remaining component, find a maximum independent set by exhaustive search. Form I as the union of these independent sets from the components.

Suppose I^* is a maximum independent set in G . By results of Al bertson and Hutchinson [Albe78], for every $\delta > 0$ all but finitely many graphs of genus g have $|I^*| \geq n(1/4 - \delta)$. Then following Lipton and Tarjan [Lipt80] we find that the relative error $(|I^*| - |I|)/|I^*|$ is $O(\sqrt{g/k(n)})$ and the time required depends on $k(n)$. For example, if we choose $k(n) = \log \log n$, the time is $O((n+g) \log n)$ and the relative error goes to 0 for fixed g as n increases.

Graphs can be embedded in nonorientable surfaces as well as orientable surfaces. There exist graphs of arbitrarily large (orientable) genus that can be embedded in the projective plane, which is the simplest nonorientable surface [Aus163]. We believe that the results of this paper will extend in a straightforward way to graphs embeddable in nonorientable surfaces. Thus we conjecture that graphs with nonorientable embeddings of Euler characteristic g satisfy a $\sqrt{gn}$ -separator theorem.

Finally, we may ask how hard it is to find the separator in Theorem 4 if we are not given an embedding of G in its genus surface. Such an embedding can be found in time polynomial in n but exponential in g [Filo79]. No embedding algorithm is known that is polynomial in the genus; indeed, several similar problems are known to be NP-complete [John82]. However, a polynomial separator algorithm would follow from a polynomial approximation algorithm that embedded the graph in a surface of genus within a constant factor of minimum.

References

[Albe78]

Michael O. Albertson and Joan P. Hutchinson.
On the independence ratio of a graph.
Journal of Graph Theory 2: 1-8, 1978.

[Batt62]

Joseph Battle, Frank Harary, Yukihiro Kodama, and J. W. T. Youngs.
Additivity of the genus of a graph.
Bulletin of the American Mathematical Society 68: 565-568, 1962.

[Aus163]

L. Auslander, T. A. Brown, and J. W. T. Youngs.
The imbedding of graphs in manifolds.
The Journal of Mathematics and Mechanics 12: 629-634, 1963.

[Djid82]

Hristo Nicolov Djidjev.
On the problem of partitioning planar graphs.
SIAM Journal on Algebraic and Discrete Methods 3: 229-240, 1982.

[Filo79]

I. S. Filotti, Gary L. Miller, and John Reif.
On determining the genus of a graph in $O(v^{O(g)})$ steps.
Proceedings of the Eleventh Annual ACM Symposium on Theory of
Computing, pages 27-37, 1979.

[Gabb79]

Ofer Gabber and Zvi Galil.
Explicit constructions of linear size superconcentrators.
Proceedings of the 20th Annual Symposium on Foundations of
Computer Science, pages 364-370, 1979.

[Gilb80]

John Russell Gilbert.
Graph Separator Theorems and Sparse Gaussian Elimination.
Ph. D. Thesis, Stanford University, 1980.

[Hoey80]

Dan Hoey and Charles E. Leiserson.
A layout for the shuffle-exchange network.
Carnegie-Mellon University Department of Computer Science
technical report CMU-CS-80-139, 1980.

[John82]

David S. Johnson.
The NP-completeness column: An ongoing guide.
Journal of Algorithms 3: 89-99, 1982.

[Leig81]

Frank Thomson Leighton.

New lower bound techniques for VLSI.

Proceedings of the 22nd Annual Symposium on Foundations of
Computer Science, pages 1-12, 1981.

[Leis80]

Charles E. Leiserson.

Area-efficient graph layouts (for VLSI).

Proceedings of the 21st Annual Symposium on Foundations of
Computer Science, pages 270-281, 1980.

[Lewi65]

P. M. Lewis II, R. E. Stearns, and J. Hartmanis.

Memory bounds for the recognition of context-free and
context-sensitive languages.

IEEE Conference Record on Switching Theory and Logical Design,
pages 191-202, 1965.

[Lipt79a]

Richard J. Lipton, Donald J. Rose, and Robert Endre Tarjan.

Generalized nested dissection.

SIAM Journal on Numerical Analysis 16: 346-358, 1979.

[Lipt80]

Richard J. Lipton and Robert Endre Tarjan.

Applications of a planar separator theorem.

SIAM Journal on Computing 9: 615-627, 1980.

[Lipt79b]

Richard J. Lipton and Robert Endre Tarjan.

A separator theorem for planar graphs.

SIAM Journal on Applied Mathematics 36: 177-189, 1979.

[Mass67]

William S. Massey.

Algebraic Topology: An Introduction.

Harcourt, Brace, and World, 1967.

[Whit73]

Arthur T. White.

Graphs, Groups, and Surfaces.

North-Holland, 1973.

Figure 1. Non-2-cell embeddings in a torus.

Figure 2. 2-cell embeddings in a torus.

Figure 3(a). H and its spanning tree embedded in a torus.

Figure 3(b).

Figure 3(c). The planarizing set H' .

Figure 4(a). A 4-pyramid.

Figure 4(b). Replacing two vertices with 4-pyramids.

Figure 5(a). A 4-pyramid. The white vertices are a file.

Figure 5(b). Two joined 4-pyramids.