Disk Packings and Planar Separators*

Daniel A. Spielman[†] U.C. Berkeley/MIT

Shang-Hua Teng[‡] University of Minnesota

Abstract

We demonstrate that the geometric separator algorithm of Miller, Teng, Thurston, and Vavasis finds a 3/4-separator of size $1.84\sqrt{n}$ for every n node planar graph. Our bound is derived from an analysis of disk packings on the sphere.

1. Introduction

Lipton and Tarjan [18] showed that every n node planar graph has a set of at most $\sqrt{8n}$ of vertices whose removal divides the rest of the graph into two disconnected pieces of size no more than $(2/3)n$. We call such a set a *2/3-separator of size* $\sqrt{8n}$. Their bound on the size of a 2/3-separator was improved to $\sqrt{6n}$ by Djidjev [7], $\sqrt{5n}$ by Gazit [14], and $\sqrt{4.5n}$ by Alon, Seymour, and Thomas [1].

Miller, Teng, Thurston, and Vavasis [22] developed a geometric approach which uses conformal mappings and centerpoints to find small separators for a class of graphs that includes planar graphs. Their technique has been improved and simplified in a series of papers [23, 21, 24, 10, 4]. A simple analysis (reviewed in Section 2) shows that their geometric separator algorithm finds a 3/4-separator of size $2\sqrt{n}$ for every n node planar graph. We improve this analysis to show that their algorithm can be used to find a 3/4-separator of size $1.84\sqrt{n}$ for every n node planar graph. Our bound is derived from an analysis of disk packings on the sphere.

In Section 2, we will review the geometric construction of Miller *et al* and give an analysis of the $2\sqrt{n}$ bound. In Section 3, we identify the slack in the analysis of the proof of Section 2, and provide a tight analysis of the how the geometric algorithm performs on average. In particular, we prove that it obtains a 3/4-separator of size at most $1.905\sqrt{n}$. This is not as good as the bound that we will eventually prove because, in many situations, the algorithm may find a separator with a better than 3/4 split. In Section 3.3, we show that the geometric algorithm can be forced to have a chance of providing a separator that exactly 3/4-splits, in which case it always produces a separator of size at most $1.84\sqrt{n}$. In contrast, Djidjev [7] constructed graphs that do not have 1/2-separators of size less than $1.65\sqrt{n}$, 1/3-separators of size less than $1.55\sqrt{n}$, or 1/4-separators of size less than $1.42\sqrt{n}$.

Since the geometric separators split with ratio 3/4 and those of [18, 14, 7, 1] split with ratio 2/3, it is not clear how one should compare the two. We try to compare them by using them in an important application—the nested dissection algorithm for Gaussian elimination of planar linear systems. In Section 4, we show that the geometric separators yield a faster nested dissection Gaussian elimination algorithm for planar linear systems than that obtained from the improvements of Lipton and Tarjan's theorem. In Section 5, we present some extensions of our planar separator theorem. In Section 6, we will discuss some potential directions for further improvement.

2. Sphere-preserving Maps and Geometric Separators

In this section, we review the geometric algorithm for finding graph separators. To provide context for the analysis of Section 3, we present a proof that this algorithm finds separators of size at most $2\sqrt{n}$.

*Errata to this paper will be available at `http://theory.lcs.mit.edu/~spielman`.

[†]Computer Science Division, U.C. Berkeley, CA 94720, `spielman@cs.berkeley.edu`. Supported in part by an NSF postdoc. After July 1996, Department of Mathematics, M.I.T., Cambridge, MA 02139

[‡]Department of Computer Science, University of Minnesota, Minneapolis, MN 55455, `steng@cs.umn.edu`. Supported in part by an NSF CAREER award (CCR-9502540) and an Intel research grant. Part of the work was done while visiting Xerox PARC in the Summer of 1995.

2.1. Definitions

Let $P = \{\boldsymbol{p}_1, ..., \boldsymbol{p}_n\}$ be a set of n points in R^d. A *centerpoint* of P is a point $\boldsymbol{c} \in \mathrm{R}^d$ such that every hyperplane passing through $\boldsymbol{c}$ divides P into two sets whose sizes have ratio at most $d : 1$. Every finite point set in R^d has a centerpoint, and one can be found using linear programming [9, Section 4]. An efficient and practical constant time centerpoint approximation algorithm can be found in [6].

Let S^d be the sphere defined by the boundary of the unit d-dimension ball. A *sphere-preserving* map from S^d to S^d is a continuous function such that the image of every sphere (of lower dimension) contained in S^d is a sphere in S^d and every sphere in S^d has a pre-image that is a sphere. Familiar sphere-preserving maps include rotations and the map that sends each point to its antipode.

We can extend the class of sphere-preserving maps over S^d to a class of "sphere-preserving" maps from R^{d-1} to S^d by first applying stereographic projection to map R^{d-1} onto S^d and then applying sphere-preserving maps on S^d [22].

We will use the following theorem concerning sphere-preserving maps of Miller *et al* [22].

Theorem 1. *Let $P = \{\boldsymbol{p}_1, ..., \boldsymbol{p}_n\}$ be a set of points in R^d. There is a sphere-preserving map Π that maps P to a set of points $Q = \Pi(P)$ on the unit sphere in R^{d+1} such that the center of the sphere is a centerpoint of Q. Such a sphere-preserving map can be found in linear time.*

We will use the following definition of graph separators:

Definition 2. A subset C of vertices of an n-vertex graph G is a δ-**separator** of size $f(n)$ if $|C| \leq f(n)$ and the vertices of $G - C$ can be partitioned into two sets A and B of size no more than δn so that there are no edges from A to B. Here, f is a function and $0 < \delta \leq 1/2$.

A graph $G = (V, E)$ is *planar* if we can "draw" it in the plane in such a way that each vertex is represented by a distinct point, each edge is represented by a continuous path between the points representing its vertices, and none of the paths representing edges intersect, except at their endpoints.

2.2. Disk Packings for Planar Graphs

The geometric separator algorithm uses the following geometric characterization of planar graphs (Koebe [17], Andreev [2, 3] and Thurston [26]): Let a *disk packing* be a set of disks $D_1, ..., D_n$ that have disjoint interiors. The *intersection graph* of the disk packing $D_1, ..., D_n$ is the graph whose vertex set is $\{D_1, ..., D_n\}$ and whose edge-set contains exactly all pairs of disks that have a common point. (See Figure 1). We will refer this type of graph as a disk packing graph.

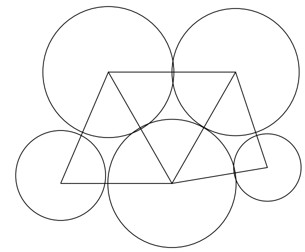

Figure 1: The intersection graph of a disk packing

It can be easily seen that every disk packing graph is a planar graph: the disk packing itself provides the straight-line planar embedding. The center of each disk represents a vertex and there is a line segment joining two centers if and only if their disks have a common point.

Theorem 3 (Koebe-Andreev-Thurston). *Every planar graph G is isomorphic to a disk packing graph.*

The Koebe-Andreev-Thurston theorem strengthens a theorem of Fáry [11] and Tutte [27, 28] that every planar graph can be embedded in the plane in such a way that each edge is mapped to a straight line segment (See [25, 12]).

2.3. Circle Separators for Planar Graphs

In this section, we provide a proof that the geometric algorithm provides a 3/4-separator of size at most $2\sqrt{n}$ for every planar graph. A similar proof can be found in Agarwal and Pach [4]. In the next section, we will examine the slack in this analysis and give a tighter analysis. We will use the following fact about convex functions.

Proposition 4. *Let $\alpha > 1$ and let $x_1, x_2, ..., x_n$ be n non-negative reals such that $\sum_{i=1}^n x_i^\alpha = 1$. Then, the function $\sum_{i=1}^n x_i$ is maximized when $x_1^\alpha = x_2^\alpha = ... = x_n^\alpha = 1/n$ [16]. Therefore, the maximum value of $\sum_{i=1}^n x_i$, subject to $\sum_{i=1}^n x_i^\alpha = 1$, is $n^{1-1/\alpha}$.*

Theorem 5. *Every n node planar graph has a 3/4-separator of size $2\sqrt{n}$.*

Proof: Let $\Gamma = \{D_1, ..., D_n\}$ be the disk packing given by Theorem 3. Let $P = \{\boldsymbol{p}_1, ..., \boldsymbol{p}_n\}$ be the centers of $\{D_1, ..., D_n\}$, respectively. Let Π be the conformal map known to exist by Theorem 1 that maps P to $Q = \Pi(P)$ on the unit sphere S_2 in R^3 so that the center of S_2 is a centerpoint of Q. Π maps Γ to a collection of caps $\Phi = \{C_1, ..., C_n\} = \{\Pi(D_1), ..., \Pi(D_n)\}$, where a *cap* on S_2 is the intersection of a halfspace with S_2. Let r_i be the radius of the circle that bounds C_i. In the following discussion, we refer to r_i as the radius of C_i. Because the center of S_2 is a center point of Q, each cap is smaller than a hemisphere.

To calculate the average number of caps of Φ that a random great circle intersects, we will use the duality between great circles and points on a unit sphere [10, 24]. Each great circle $\mathcal{G}$ can be identified with the pair of points $p_{\mathcal{G}}$ and $q_{\mathcal{G}}$ on S_2 that lie on the axis of $\mathcal{G}$. We call $\mathcal{G}$ the dual of $p_{\mathcal{G}}$ and $q_{\mathcal{G}}$. For each pair of great circles $\mathcal{G}$ and $\mathcal{G}'$ of S_2, $\mathcal{G}$ contains $p_{\mathcal{G}'}$ (and hence $q_{\mathcal{G}'}$ as well) if and only if $\mathcal{G}'$ contains $p_{\mathcal{G}}$ (and hence $q_{\mathcal{G}}$).

Define a *great belt* to be a set of points of S_2 that lie between a pair of parallel planes symmetric about the center of S_2. The *width* of a great belt is then the distance between its two planes. Notice that a great circle is a great belt with width 0. As observed in [10, 24], the dual of each cap (the union of all great circles dual to points in the cap) is a great belt. Here we assume that each cap is smaller than a hemisphere. Moreover, if the radius of a cap is r, then the width of its great belt is $2r$.

Let $B_1, ..., B_n$ be the great belts that are the duals of the caps $C_1, ..., C_n$, respectively. The point-great-circle duality extends to caps and their dual great belts, as observed in [10, 24]: For each cap C_i and great circle $\mathcal{G}$, $\mathcal{G}$ intersects C_i if and only if the dual points of $\mathcal{G}$, $p_{\mathcal{G}}$ and $q_{\mathcal{G}}$, are in B_i. Therefore, the average number of caps that a random great circle intersects, $\mathrm{Avg}(\Phi)$, is equal to the total area of the great belts $B_1, ..., B_n$ divided by the surface area of S_2,

$$\mathrm{Avg}(\Phi) = \sum_{i=1}^{n} \mathrm{Area}(B_i)/4\pi = \sum_{i=1}^{n} r_i. \qquad (1)$$

The last equality follows from the fact[1] that $\mathrm{Area}(B_i) = 4\pi r_i$.

Because Γ is a disk packing, we know that $\sum_{i=1}^{n} \mathrm{Area}(C_i) \leq 4\pi$, where 4π is equal to the surface area of a unit sphere in R^3. Combining this fact with $\mathrm{Area}(C_i) \geq \pi r_i^2$, we find

$$\sum_{i=1}^{n} \pi r_i^2 \leq \sum_{i=1}^{n} \mathrm{Area}(C_i) \leq 4\pi,$$

implying

$$\sum_{i=1}^{n} r_i^2 \leq 4. \qquad (2)$$

Applying Proposition 4 to Equations 2 and 1 we learn that

$$\mathrm{Avg}(\Phi) \leq 2\sqrt{n}.$$

□

[1]We would like to thank Dafna Talmor of CMU for pointing this fact out to us.

3. Better Separator Bounds

The proof of Theorem 5 seems to indicate that uniform disk packings on the unit sphere in R^3 are the worst case for the geometric algorithm. If uniform disk packings are indeed the worst case, then we should get a better separator bound, because uniform disks cannot completely cover the sphere and hence the total area of the caps would be less than 4π. However, the proof of Theorem 5 did not quite prove that uniform disk packings are the worst case because non-uniform disks can cover the sphere better. On the other hand, a uniform disk packing has the best ratio of cap area to great belt area. In fact, the Djidjev [7] examined the graphs obtained from regular disk packings and showed that they do not have 1/3-separators of size less than $1.55\sqrt{n}$. One can use similar techinques to show that they do not have 1/2-separators of size less than $1.65\sqrt{n}$ or 1/4-separators of size less than $1.42\sqrt{n}$.

In this section, we show that a uniform packing of disks maximizes the expected number of caps intersected by a random great circle. We can incorporate this fact into the proof of Theorem 5 to get a better bound on the size of the separator provided by the geometric algorithm.

We will eventually prove:

Theorem 6. *Every n node planar graph has a 3/4-separator of size $1.84\sqrt{n}$.*

3.1. Polytopes of Disk Packings

Let $\{C_1, \ldots, C_n\}$ be a packing of caps on the sphere. To each cap C_i, we associate the plane P_i such that the perimeter of C_i is the intersection of P_i with the sphere. We will occasionally abuse C_i by using it to refer to the circle defined by the intersection of the cap with the plane. To each plane P_i, we associate the half-space H_i containing the origin whose boundary is P_i.[2] To the packing of caps, we will associate the polytope $\mathcal{P}(C_1, \ldots, C_n)$ defined by the intersections of the H_i's.[3]

Later in our proof, we will want the surface of this polytope to be close to the surface of the sphere. Since this is not necessarily the case for all packings of caps on the sphere, we will show that one can add a few extra disks to any packing of caps on the sphere so that the resulting polytope is close to the sphere, provided that the original caps were not too big.

By the *radius* of a cap, we mean the radius of the

[2]If we assume that the cap has radius less than one, then this half-space is uniquely defined.

[3]This polytope could be unbounded, but we will make sure that this does not happen.

circle formed by the intersection of the cap with its associated plane. So that we can describe caps from the perspective of the sphere on which they lie, we will say that a cap *occupies α radians* if α is the angle at the tip of the cone emanating from the center of the sphere and passing through the circle associated with the cap. Thus, a cap that occupies α radians on the sphere has radius $\sin(\alpha/2)$. We will say that a line segment occupies α radians if the triangle determined by the line segment and the center of the sphere has angle α at the center of the sphere.

Lemma 7. *Let $\epsilon > 0$ and let $\{C_1, \ldots, C_n\}$ be a packing of caps on the sphere in which each cap occupies at most 2ϵ radians. Then, there exist $e = O(1/\epsilon^2)$ caps $E_1, \ldots, E_e$ that occupy 2ϵ radians so that every point on the surface of the polytope $\mathcal{P}(C_1, \ldots, C_n, E_1, \ldots, E_e)$ is at distance at most $\sin(5\epsilon)$ from the sphere.*

Proof: We choose the caps $E_1, \ldots, E_e$ by brute force: if there is a gap in the packing into which we can insert a cap of 2ϵ radians, then insert the cap and add it to our list. Because each such cap has area $\theta(\epsilon^2)$, we can insert at most $O(1/\epsilon^2)$ of them.

We now wish to show that, after these caps have been inserted, no point on any polygon can be at distance more than $\sin(5\epsilon)$ from the surface of the sphere. Consider a point on one of the polygons that bound the polytope (see Figure 2). We can draw a line l in the plane of the polygon from the point to the sphere. We draw this line so that it is perpendicular to the circle that bounds the cap associated with the plane.

Let β be the number of radians that the line l occupies, and let q be the point where the line through p and the center of the sphere intersects the surface of the sphere (see Figure 2). The spherical cap occupying 2β radians centered at q cannot contain the center of any cap in the packing. If this cap contained the center point of any cap in the packing, then the plane associated with that cap would separate p from the center of the sphere, which would be a contradiction. Thus, $\beta < 4\epsilon$, or else it would be possible to insert another cap of radius ϵ into the packing.

Because $\beta < 4\epsilon$ and the cap contained in the polygon occupies at most 2ϵ radians, we know that the length of l is at most $\sin(5\epsilon)$. $\square$

3.2. A Packing Theorem

Theorem 8. *Let $\{C_1, \ldots, C_n\}$ be a packing of caps on the sphere of radii $r_1, \ldots, r_n$ respectively. Then,*

$$\sum_{i=1}^{n} r_i \leq \sqrt{\frac{2\pi}{\sqrt{3}}}(1 + o(1))\sqrt{n}.$$

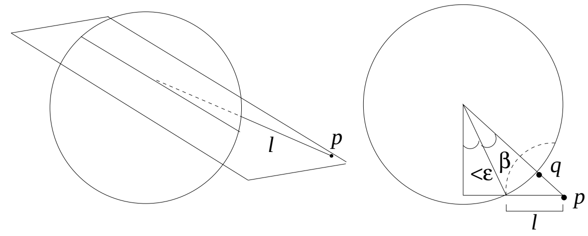

Figure 2: a point on a polygon

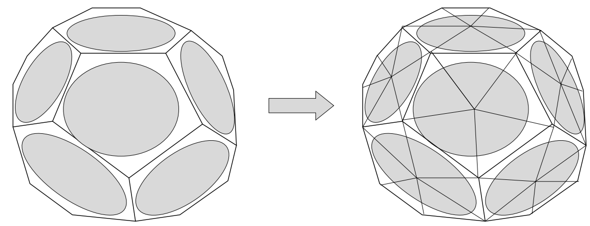

Figure 3: Triangulating each polygon.

Accordingly, the geometric separator algorithm always finds a separator of size at most $1.905\sqrt{n}$.

The bound of Theorem 8 is tight and can be achieved by uniform disk packings.

Proof: We begin by removing all the big caps from the packing. Choose an $\epsilon > 0$, and let $\{C_{m+1}, \ldots, C_n\}$ be the caps in the packing of radius greater than ϵ. There are at most $O(1/\epsilon^2)$ such caps, and these contribute at most $O(1/\epsilon^2)$ to the sum of the radii. We now consider the packing $\{C_1, \ldots, C_m\}$

Strangely enough, we will now throw a collection of caps of radius ϵ back in to the packing. Let $E_1, \ldots, E_e$ be a maximal collection of caps of radius ϵ such that $\{C_1, \ldots, C_m, E_1, \ldots, E_e\}$ is a packing. Clearly, $e = O(1/\epsilon^2)$. For convenience of notation, let $p = m + e$, let $\{C'_1, \ldots, C'_p\} = \{C_1, \ldots, C_m, E_1, \ldots, E_e\}$, and let r'_i be the radius of cap C'_i. Consider the polytope $\mathcal{P}(C'_1, \ldots, C'_p)$. By Lemma 7, this polytope lies within the sphere of radius $1+\sin(5\epsilon)$. Thus, the surface area of the polytope is at most $4\pi(1 + \sin(5\epsilon))$.

Recall that the surface of the polytope is composed of polygons, each of which contains a circle corresponding to the perimeter of one of the caps. Triangulate each such polygon by adding edges connecting the center of the circle it contains to each vertex of the polygon (see Figure 3). We will now prove the theorem by counting the contribution of each triangle.

For circle C'_i, let the triangles surrounding it be $t_{i,1}, \ldots, t_{i,n_i}$ and let $\alpha_{i,j}$ be the angle formed by $t_{i,j}$ at the center of the circle. By Lemma 9,

$$\text{Area}(t_{i,j}) \geq r_i^2 \tan(\alpha_{i,j}/2).$$

We now consider the maximum of

$$\sum_{i=1}^{p} r_i = \frac{1}{2\pi}\sum_{i=1}^{p}\sum_{j=1}^{n_i} r_i \alpha_{i,j}$$

subject to

$$\sum_{i=1}^{p}\sum_{j=1}^{n_i} r_i^2 \tan(\alpha_{i,j}/2) \leq 4\pi(1+\sin(5\epsilon)) \qquad \text{and}$$

$$\sum_{j=1}^{n_i} \alpha_{i,j} = 2\pi, \text{ for all } i.$$

Weaken the constraints to obtain a maximum that is at least as large by letting triangle $t_{i,j}$ have area $r_{i,j}^2 \tan \alpha_{i,j}$ and taking the maximum of

$$\frac{1}{2\pi}\sum_{i=1}^{p}\sum_{j=1}^{n_i} r_{i,j}\alpha_{i,j}$$

subject to

$$\sum_{i=1}^{p}\sum_{j=1}^{n_i} r_{i,j}^2 \tan(\alpha_{i,j}/2) \leq 4\pi(1+\sin(5\epsilon)) \qquad \text{and}$$

$$\sum_{i=1}^{p}\sum_{j=1}^{n_i} \alpha_{i,j} \leq 2\pi p.$$

By Lemma 10, this maximum is obtained when all the $r_{i,j}$'s are equal to some value r and all the $\alpha_{i,j}$'s are equal to some value α. Let q be the number of triangles in our collection. Because the edge-graph of the polytope is planar, $q < 6p$. Thus, we have $\alpha = 2\pi p/q \geq \pi/3$ and

$$q(r^2 \tan \pi/6) \leq 4\pi(1+\sin 5\epsilon), \text{ so}$$

$$r \leq \sqrt{\frac{4\pi(1+\sin 5\epsilon)}{q \tan \pi/6}}.$$

Because $q < 6p$, we find that

$$\sum_{i=1}^{p} r_i \;\leq\; p\sqrt{\frac{4\pi(1+\sin 5\epsilon)}{6p\tan\pi/6}} \;\leq\; \sqrt{\frac{2\pi}{\sqrt{3}}}\sqrt{1+\sin 5\epsilon}\sqrt{p}$$

As we let ϵ go to zero, we find that

$$\sum_{i=1}^{n} r_i \;\leq\; \sum_{i=1}^{p} r'_i + O(1) \;\leq\; \sqrt{\frac{2\pi}{\sqrt{3}}}(1+o(1))\sqrt{p} \;\leq\; 1.905\sqrt{n}.$$

□

We now prove the remaining lemmas that we used in the proof of Theorem 1.

Lemma 9. *Let C be a circle of radius r and let T be a triangle with vertices a, b, and c such that a is the center of C and edge $\overline{bc}$ is contained in the exterior of C. Let α be the angle of the triangle at a (see Figure 4(a)). Then, the area of T is at least $r^2\tan(\alpha/2)$.*

Proof: If we move the edge $\overline{bc}$, without changing its slope, until it touches C, then we obtain a similar triangle with area at most the area of T (see Figure 4(b)). Thus, we can assume that edge $\overline{bc}$ touches C.

Draw a line from the center of C to the point at which edge $\overline{bc}$ touches C. This line divides T into two triangles, with angles that we will call β and γ (see Figure 4(c)). We can now compute the area of T to be

$$\text{Area}(T) = \frac{r^2(\tan\beta + \tan\gamma)}{2}.$$

Because β and γ are less than $\pi/2$ and $\tan(x)$ is convex for $x \in [0, \pi/2]$, we can show that T will have minimal area if $\beta = \gamma$, in which case it has area $r^2\tan(\alpha/2)$. □

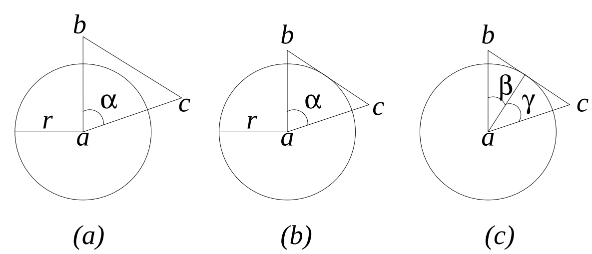

Figure 4:

Lemma 10. *Let $f(x) = x^2$ and $g(y) = \tan(y)$. Then, for variables $y_i \in [0, \pi/2]$ and $x_i \geq 0$, the maximum of*

$$\sum_{i=1}^{n} x_i y_i \qquad (*)$$

subject to

$$\sum f(x_i)g(y_i) \;\leq\; c_1, \qquad \textit{and} \qquad (3)$$

$$\sum y_i \;\leq\; c_2, \qquad (4)$$

for $c_1, c_2 > 0$, is achieved when all the x_i's are equal and all the y_i's are equal.

We will use the follwing lemma in the proof of Lemma 10.

Lemma 11. *Let $a > 0$ and let $f(x)$ and $g(y)$ be functions such that*

(a) *$f(x)$ is a strictly increasing convex function from the non-negative reals to the non-negative reals,*

(b) *$g(y)$ is a strictly increasing convex function from the interval $[0, a]$ to the non-negative reals.*

Then, for variables $a \geq y_1 \geq y_2 \geq \cdots \geq y_n \geq 0$, and $x_i \geq 0$, the maximum of

$$\sum_{i=1}^{n} x_i y_i \qquad (*)$$

subject to

$$\sum f(x_i) g(y_i) \leq c,$$

is achieved when $0 \leq x_1 \leq x_2 \leq \cdots \leq x_n$.

Proof: Let $x_1, \ldots, x_n$ and $y_1, \ldots, y_n$ be the values that maximize $(*)$ subject to the constraint. Assume, by way of contradiction, that $x_i > x_j$ for $i < j$. We will show that it is possible to adjust x_i and x_j so that $(*)$ is unchanged but the constraint has slack. We could then slightly increase all the x_i's without violating the constraint. As this would increase $(*)$, we would have a contradiction.

Let

$$x_i' = x_i - \epsilon y_j, \qquad \text{and} \qquad x_j' = x_j + \epsilon y_i.$$

If we replace x_i with x_i' and x_j with x_j', $(*)$ is unaffected; but, we will see that slack is introduced in the constraint. As we take the limit as ϵ goes to zero, we compute

$$f(x_i)g(y_i) + f(y_j)g(y_j) - f(x_i')g(y_i) - f(y_j')g(y_j) \\ \rightarrow \epsilon(y_j f'(x_i) g(y_i) - y_i f'(x_j) g(y_j)).$$

Because $x_i > x_j$, we know that $f'(x_i) > f'(x_j)$ by property (a). Thus, slack will be introduced if

$$y_i g(y_j) \leq y_j g(y_i) \Leftrightarrow \frac{g(y_i)}{y_i} \geq \frac{g(y_j)}{y_j},$$

which we know because property (b) implies that $g(y)/y$ is strictly increasing and $y_i \geq y_j$. □

Proof: [**of Lemma 10**] Let $x_1, \ldots, x_n$ and $y_1, \ldots, y_n$ be the values that maximize $(*)$ subject to the constraints. By Lemma 11, we can assume that $x_1 \leq x_2 \leq \cdots \leq x_n$ and $y_1 \geq y_2 \geq \cdots \geq y_n$. Let $i < j$ and consider what happens if we replace y_i with y_i' and y_j with y_j' where

$$y_i' = y_i - \epsilon x_j, \qquad \text{and} \qquad y_j' = y_j + \epsilon x_i,$$

for $\epsilon > 0$. The sum $(*)$ is unaffected, and constraint 4 will not be violated because $x_j \geq x_i$. As ϵ approaches zero from above, constraint 3 will develop slack if

$$\epsilon(x_j g'(y_i) f(x_i) - x_i g'(y_j) f(x_j)) > 0.$$

This will happen unless

$$\begin{aligned} x_j g'(y_i) f(x_i) &\leq x_i g'(y_j) f(x_j) \\ \frac{f(x_i) x_j}{f(x_j) x_i} &\leq \frac{g'(y_j)}{g'(y_i)} \\ \frac{x_i}{x_j} &\leq \frac{\cos^2(y_i)}{\cos^2(y_j)} \end{aligned}$$

Similarly, if we consider the substitution

$$x_i' = x_i + \delta y_j, \qquad \text{and} \qquad x_j' = x_j - \delta y_i,$$

for $\delta > 0$ and small, then sum $(*)$ will remain unchanged and constraint 3 will develop slack if

$$y_i f'(x_j) g(y_j) - y_j f'(x_i) g(y_i) > 0.$$

Since we assume that we cannot obtain such slack, we find that

$$\begin{aligned} y_j f'(x_i) g(y_i) &\geq y_i f'(x_j) g(y_j) \\ \frac{f'(x_i)}{f'(x_j)} &\geq \frac{y_i g(y_j)}{y_j g(y_i)} \\ \frac{x_i}{x_j} &\geq \frac{y_i \tan(y_j)}{y_j \tan(y_i)}. \end{aligned}$$

By combining this inequality with the one we developed before, we find that slack can be introduced in constraint 3 unless

$$\begin{aligned} \frac{\cos^2(y_i)}{\cos^2(y_j)} &\geq \frac{y_i \tan(y_j)}{y_j \tan(y_i)} \\ \frac{y_j}{\sin(y_j)\cos(y_j)} &\geq \frac{y_i}{\sin(y_i)\cos(y_i)} \end{aligned}$$

However, this can only occur if $y_j = y_i$, because the function $\frac{y}{\sin(y)\cos(y)} = \frac{2y}{\sin(2y)}$ is strictly increasing for $0 \leq y \leq \pi/2$. Thus, all of the y_i's must be equal. By again combining the same two inequalities,

$$1 = \frac{\cos^2(y_i)}{\cos^2(y_j)} \geq \frac{x_i}{x_j} \geq \frac{y_i \tan(y_j)}{y_j \tan(y_i)} = 1,$$

we find that all the x_i's must be equal as well. □

3.3. If the Center is a δ-Centerpoint

The separator size bound of Theorem 5 is not tight for 3/4-separators because, in the uniform disk packings that produce the largest separators, every plane through the the center of the sphere splits the caps into two sets of almost the same size. In other words, the center of the sphere is a much stronger center point than one needs to obtain a 3/4-separator.

From the linear programming characterization of centerpoints [9], one can see that every point set in d dimensions has a centerpoint and a hyperplane passing through it that can be used to split the points with ratio $1 : d$. If we map such a centerpoint to the center of the sphere, then we can show that a cut by a random hyperplane through the center provides a smaller separator than that guaranteed by Theorem 8.

Let $\mathcal{H}$ be the hyperplane through the center of the sphere such that one half-space defined by the hyperplane contains the centers of at most $n/4$ of the caps in Φ. Using the assumption that no cap in Φ has radius greater than ϵ, we can use arguments similar to those in the proof of Theorem 8 to show that the sum of the radii of the caps will be maximized when all of the caps in each half-space have the same radius. In this case, we expect that the geometric separator algorithm will return a separator of size at most

$$\mathrm{Avg}(\Phi) \leq \sqrt{\frac{2\pi}{\sqrt{3}}}\left(\frac{1+\sqrt{3}}{2\sqrt{2}} + o(1)\right)\sqrt{n},$$

which is less than $1.84\sqrt{n}$.

Similarly, for any $k \geq 3$, we can establish the following bound.

Corollary 12. *For any $k \geq 3$, every planar graph has a $k/(k+1)$-separator of size*

$$\sqrt{\frac{2\pi}{\sqrt{3}}}\left(\frac{(1+\sqrt{k})}{\sqrt{2(1+k)}} + o(1)\right)\sqrt{n}.$$

4. Quality Measurement by Nested Dissection

It is not immediately apparent how we should compare our 3/4-separators of size $1.905\sqrt{n}$ with the 2/3-separators of size $\sqrt{4.5n}$ obtained by Alon, Seymour, and Thomas. In this section, we will compare these separators by examining how their use affects one of the most important applications of small separator algorithms—the *nested dissection* algorithm for solving sparse linear systems [13, 19, 22].

The nested dissection algorithm solves a system of linear equations by finding a small *separator* of the graph of the system, eliminating the variables in the two subgraphs created by the separator, and then eliminating the variables corresponding to vertices in the separator. For a class of graphs with a family of δ-separators of size $f(n)$ ($0 < \delta \leq 1/2$), the worst case time complexity of the nested dissection algorithm is given by the recurrence

$$T(n) = T(\delta n) + T((1-\delta)n) + (f(n))^3, \qquad (5)$$

where $T(n)$ denotes the worst case time complexity for solving an n variable linear system over the class of graphs.

Using the planar separator theorem of Lipton and Tarjan [18], Lipton, Rose, and Tarjan [19] showed that every planar linear system of n variables can be solved in $O(n^{1.5})$ time. We compare our separators with those obtained by Alon, Seymour, and Thomas by comparing the constants that appear in front of the $n^{1.5}$ term.

Lemma 13. *The worst case time complexity, $T_{\alpha,\delta}(n)$, of nested dissection using δ-separators of size $\alpha\sqrt{n}$ is*

$$\frac{\alpha^3 n^{1.5}}{1-\delta^{1.5}-(1-\delta)^{1.5}}.$$

Proof: By Equation 5, we have

$$T_{\alpha,\delta}(n) = T_{\alpha,\delta}(\delta n) + T_{\alpha,\delta}((1-\delta)n) + \alpha^3 n^{1.5}.$$

Let the solution of $T_{\alpha,\delta}(n)$ be $\alpha^3 c n^{1.5}$, we have

$$\alpha^3 c n^{1.5} = \alpha^3 c \delta^{1.5} n^{1.5} + \alpha^3 c (1-\delta)^{1.5} n^{1.5} + \alpha^3 n^{1.5},$$

and thus $c = 1/(1-\delta^{1.5}-(1-\delta)^{1.5})$. □

From Alon-Seymour-Thomas's 2/3-separators of size $\sqrt{4.5n}$ we obtain

$$T_{\sqrt{4.5},2/3} = 36.2662n^{1.5},$$

whereas, from our 3/4-separators of size 1.84, we obtain

$$T_{1.84,3/4} = 27.6276n^{1.5}.$$

In a recent conversion, Djidjev [8] indicated that he can prove that every n node planar graph has a 2/3-separator of size $2\sqrt{n}$, from which one would obtain

$$T_{2,2/3} = 30.3930n^{1.5}.$$

5. Extensions

Gazit and Miller [15] proved the following interesting edge-separator theorem for planar graphs.

Theorem 14 (Gazit–Miller). *If G is a planar graph of n nodes with degrees $d_1, \ldots, d_n$, then G has a 1/3-edge separator of size $1.58\sqrt{\sum_{i=1}^{n} d_i^2}$.*

They prove this result by carefully analyzing the simple-cycle separator construction of Miller [20]. We now give a simple proof of a similar statement.

Lemma 15. *Every planar graph G of n nodes has a $(1/4 - o(1))$-edge separator of size*

$$\sqrt{\sum_{i=1}^{n} d_i^2}.$$

Proof: The simplest way to convert a vertex-separator into an edge-separator is to remove all edges incident to vertices in the separator and to divide these vertices among the two subgraphs. The resulting edge separator has size bounded by the total number of edges incident to the vertex separator. We can reduce this number by a factor of 2 by assigning the all vertices in the vertex separator to one of the subgraphs. From a δ-separator of size $o(n)$, we obtain a $(\delta - o(1))$-edge separator. Thus, a random great circle (as in the proof of Theorem 6) induces an edge-separator of expected size

$$h(\vec{r}) = \frac{1}{2}\sum_{i=1}^{n} d_i r_i,$$

where r_i is the radius of the ith cap after we map a Koebe embedding of G onto the unit sphere as in Theorem 1, and $\vec{r} = (r_1, ..., r_n)$. Again, $\vec{r}$ satisfies

$$g(\vec{r}) = \sum_{i=1}^{n} r_i^2 \leq 4.$$

We can use Lagrange's method to find the maximum of $h(\vec{r})$ subject to $g(\vec{r}) \leq 4$. We will overestimate the separator size if we allow $g(\vec{r}) = 4$. Define $f(\vec{r}, \lambda) = h(\vec{r}) - \lambda(g(\vec{r}) - 4)$. $h(\vec{r})$ is maximized when $(\partial f)/(\partial r_i) = d_i/2 - 2\lambda r_i = 0$ for all i, implying $\lambda = d_1/(4r_1) = d_2/(4r_2) = ... = d_n/(4r_n)$. Thus, $r_i = d_i/(4\lambda)$. Combining this with $\sum_{i=1}^{n} r_i^2 = 4$, we find $\lambda = \sqrt{\sum_{i=1}^{n} d_i^2}/8$, which implies $r_i = 2d_i/\sqrt{\sum_{i=1}^{n} d_i^2}$. Therefore,

$$h(\vec{r}) = \frac{1}{2}\sum_{i=1}^{n} d_i r_i \leq \sqrt{\sum_{i=1}^{n} d_i^2}.$$

□

We can extend this edge-separator result to intersection graphs [21] in higher dimensions. Intersection graphs are defined by neighborhood systems [21]. A *neighborhood system* is a set of closed balls in Euclidean space. A *k-ply* neighborhood system is one in which no point is contained in the interior of more than k of the balls. Given a neighborhood system, $\Gamma = \{B_1, \ldots, B_n\}$, we define the *intersection graph* of Γ to be the undirected graph with vertex set $V = \{B_1, \ldots, B_n\}$ and edge set

$$E = \{(B_i, B_j) : B_i \cap B_j \neq \emptyset\}.$$

Lemma 16. *Let G be the intersection graph of a k-ply neighborhood system in R^d. If G has n nodes of degrees $d_1, \ldots, d_n$, then G has a $1/(k+2)$-edge separator of size*

$$O\left(\sum_{i=1}^{n} d_i^{d/(d-1)}\right)^{1-1/d},$$

where d_i is the degree of the ith vertex in G.

Proof: As in the proof of Proposition 15, we can bound the expected size of the edge separator by

$$\max \sum_{i=1}^{n} d_i r_i,$$

subject to

$$\sum_{i=1}^{n} r_i^d = O(1).$$

By Lagrange's method, we can bound this by

$$O\left(\sum_{i=1}^{n} d_i^{d/(d-1)}\right)^{1-1/d}.$$

□

The argument of Section 2 can be easily extended to higher dimensions to obtain the following improvement of a lemma from Miller *et al.*

Lemma 17. *Let $\Phi = \{C_1, ..., C_n\}$ be a set of caps on the unit d-sphere in R^{d+1}. Let $k = (\sum_{i=1}^{n} \mathrm{Area}(C_i))/A_d$. Let $\mathrm{Avg}(\Phi)$ be the average number of caps that a random great sphere intersects. Then*

$$\mathrm{Avg}(\Phi) \leq \left(\frac{2A_{d-1}}{A_d^{1-1/d} V_d^{1/d}} + o(1)\right) k^{1/d} n^{1-1/d},$$

where V_d is the volume of a unit d-dimensional ball and A_d is the surface area of a unit d-sphere in R^{d+1}. We have $V_d = (2\pi^{d/2})/(d\Gamma(d/2))$, where $\Gamma(x)$ is the *gamma-function*, and $A_{d-1} = dV_d$.

The convex function argument also implies that the worst case is achieved when all caps are of equal size. Lemma 17 is stronger than the original presented in Miller *et al* in that:

i. The original assumed that the caps $\Phi = \{C_1, ..., C_n\}$ were k-ply in the sense that there was no point on the sphere strictly interior to more than k caps. The improved lemma has been applied by Cao, Gilbert, and Teng [5] in a new hyperplane-based separator theorem.

ii. The constant factor in the original was much larger and was not explicitly given.

The proof given in Section 2 is much simpler and provides a better constant than the original proof of Miller *et al* [22]. It also provides a better understanding of on which graphs the geometric algorithm provides the largest separators.

6. Conclusions and Open Questions

It remains to prove tight upper and lower bounds on the size of a δ-separator of a planar graph for any $\delta > 0$.

We conclude by observing that it should be possible to improve the argument in Section 3.3. It seems to us that, unless the distribution of caps on the sphere is highly uneven, a random hyperplane cut will provide a cut of ratio better than $1 : 3$. In those cases where the distribution is sufficiently uneven that we do expect a ratio of $1 : 3$, we expect that the size of the separator will be smaller than we have indicated.

7. Acknowledgements

We would like to thank Jeff Erickson, John Gilbert, Gary Miller, and Dafna Talmor for helpful discussions of the material in this paper.

References

[1] N. Alon, P. Seymour, and R. Thomas. Planar separators, *SIAM J. Disc. Math.*, 2(7): 184–193, May, 1994.

[2] E. M. Andreev. On convex polyhedra in Lobacevskii space. *Math. USSR Sbornik*, 10(3):413–440, 1970.

[3] E. M. Andreev. On convex polyhedra of finite volume in Lobacevskii space. *Math. USSR Sbornik*, 12(2):270–259, 1970.

[4] J. Pach and P. Agarwal. *Combinatorial Geometry.* Wiley-Interscience, 1995.

[5] F. Cao, J. R. Gilbert, and S.-H. Teng. Partitioning Unstructured Meshes with Lines and Planes. *the 9th International Conference on Domain Decomposition.* submitted, 1996.

[6] K. Clarkson, D. Eppstein, G. L. Miller, C. Sturtivant, and S.-H. Teng. Approximating center points with and without linear programming. In *Proceedings of 9th ACM Symposium on Computational Geometry*, pages 91–98, 1993.

[7] H. N. Djidjev. On the problem of partitioning planar graphs. *SIAM J. Alg. Disc. Math*, 3(2): 229-240, June, 1982.

[8] H. N. Djidjev. Personal communication, 1995, citing H. N. Djidjev and S. M. Venkatesan, Reduced constants for simple cycle graph separation, *Acta Informatica* (to appear).

[9] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, NY, 1987.

[10] D. Eppstein, G. L. Miller, and S.-H. Teng. A deterministic linear time algorithm for geometric separators and its applications. In *Proceedings of 9th ACM Symposium on Computational Geometry*, pages 99–108, 1993.

[11] I. Fáry. On straight line representation of planar graphs. *Acta. Sci. Math.* 11: 229–233, 1948.

[12] H. De Fraysseix, J. Pach, and R. Pollack. Small sets supporting Fáry embeddings of planar graphs. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, 426–433, 1988.

[13] J. A. George and J. W. H. Liu. *Computer Solution of Large Sparse Positive Definite Systems.* Prentice-Hall, 1981.

[14] H. Gazit. An Improved Algorithm for Separating a Planar Graph, manuscript, USC, 1986.

[15] H. Gazit and G. L. Miller. Planar separators and the Euclidean norm. *SIGAL International Symposium on Algorithms*, Information Processing Society of Japan, Tokyo, August, 1990, Springer-Verlag.

[16] G. Hardy, J. E. Littlewood, and G. Pólya. *Inequalities.* Cambridge University Press, second edition, 1952.

[17] P. Koebe. Kontaktprobleme der konformen Abbildung. *Ber. Verh. Sächs. Akademie der Wissenschaften Leipzig, Math.-Phys. Klasse*, 88:141–164, 1936.

[18] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM J. of Appl. Math.*, 36:177–189, April 1979.

[19] R. J. Lipton, D. J. Rose, and R. E. Tarjan. Generalized nested dissection. *SIAM J. on Numerical Analysis*, 16:346–358, 1979.

[20] G. L. Miller. Finding small simple cycle separators for 2-connected planar graphs. *Journal of Computer and System Sciences*, 32(3):265–279, June 1986.

[21] G. L. Miller, S.-H. Teng, and S. A. Vavasis. An unified geometric approach to graph separators. In *32nd Annual Symposium on Foundations of Computer Science*, pages 538–547. IEEE, 1991.

[22] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. Geometric separators for finite element meshes. Submitted to *SIAM J. Scientific Computing*, 1995.

[23] G. L Miller and W. Thurston. Separators in two and three dimensions. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, pages 300–309, Maryland, May 1990. ACM.

[24] S.-H. Teng. *Points, Spheres, and Separators: a unified geometric approach to graph partitioning*. PhD thesis, Carnegie-Mellon University, School of Computer Science, 1991. CMU-CS-91-184.

[25] C. Thomassen. Planarity and duality of finite and infinite graphs. *Journal of Combinatorial Theory*, Series B, 29: 244–271, 1980.

[26] W. P. Thurston. *The geometry and topology of 3-manifolds*. Princeton University Notes, 1988.

[27] W. T. Tutte. Convex representations of graphs. *Proc. London Math. Soc.* 10(3): 304–320, 1960.

[28] W. T. Tutte. How to draw a graph. *Proc. London Math. Soc.* 13(3): 743–768, 1963.