Separators for Sphere-Packings and Nearest Neighbor Graphs

GARY L. MILLER

Carnegie Mellon University, Pittsburgh, Pennsylvania

SHANG-HUA TENG

University of Minnesota, Minneapolis, Minnesota

WILLIAM THURSTON

University of California, Berkeley, California

AND

STEPHEN A. VAVASIS

Cornell University, Ithaca, New York

Abstract. A collection of n balls in d dimensions forms a k-ply system if no point in the space is covered by more than k balls. We show that for every k-ply system Γ, there is a sphere S that intersects at most $O(k^{1/d}n^{1-1/d})$ balls of Γ and divides the remainder of Γ into two parts: those in the interior and those in the exterior of the sphere S, respectively, so that the larger part contains at most $(1 - 1/(d + 2))n$ balls. This bound of $O(k^{1/d}n^{1-1/d})$ is the best possible in both n and k. We also present a simple randomized algorithm to find such a sphere in $O(n)$ time. Our result implies that every k-nearest neighbor graphs of n points in d dimensions has a separator of size $O(k^{1/d}n^{1-1/d})$. In conjunction with a result of Koebe that every triangulated planar graph is isomorphic to the intersection graph of a disk-packing, our result not only gives a new geometric proof of the planar separator theorem of Lipton and Tarjan, but also generalizes it to higher dimensions. The separator algorithm can be used for point location and geometric divide and conquer in a fixed dimensional space.

G. L. Miller was supported in part by National Science Foundation grant CCR 90-16641.

S.-H. Teng was supported by an NSF CAREER award (CCR 95-02540), an Alfred P. Sloan Research Fellowship, and an Intel research grant. Part of the work was done while S.-H. Teng was at MIT, Xerox Corporation (PARC), and Carnegie Mellon University.

S. A. Vavasis was supported by an NSF Presidential Young Investigator award.

Authors' addresses: G. L. Miller, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213; S.-H. Teng, Department of Computer Science, University of Minnesota, Minneapolis, MN 55455; W. Thurston, Department of Mathematics, University of California, Berkeley, CA 94720; S. A. Vavasis, Department of Computer Science, Cornell University, Ithaca, NY 14853.

Categories and Subject Descriptors: [**Data Structures**]: *graphs*, trees; F.2.1 [**Analysis of Algorithms and Problem Complexity**]: Numerical Algorithms and Problems—*computation of transforms (e.g., Fast Fourier transform); computations on matrices*; F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*computations on discrete structures; geometrical problems and computations; sorting and searching*; G.2.1 [**Discrete Mathematics**]: Combinatorics; G.2.2 [**Discrete Mathematics**]: Graph Theory—*graph algorithms; trees*; G.3 [**Probability and Statistics**]: *probabilistic algorithms (including Monte Carlo); random number generation*; G.4 [**Mathematical Software**]: *algorithm analysis; efficiency*

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Centerpoint, computational geometry, graph algorithms, graph separators, partitioning, probabilistic method, point location, randomized algorithm, sphere-preserving mapping

1. *Introduction*

Motivations of this work are the planar separator theorem of Lipton and Tarjan [1979], the geometric characterization of planar graphs of Koebe [1936] (see also Andreev [1970a; 1970b] and Thurston [1988]), and geometric divide and conquer.

In 1979, Lipton and Tarjan [1979] gave a linear time algorithm that divides any n-vertex planar graph into two disconnected subgraphs each has size no more than $(2/3)n$ by removing at most $\sqrt{8n}$ vertices. Their result improved a theorem of Ungar [1951] who showed it is sufficient to remove $O(\sqrt{n}\,\log n)$ vertices to partition a planar graph. A subset of vertices whose removal divides a graph into two subgraphs of roughly equal size, as given in the results above, is called a *separator* of the graph (see Section 3 for the definition).

Separators are most useful for designing efficient divide and conquer graph algorithms. The planar separator theorem of Lipton and Tarjan has been used in the solution of planar linear systems [Lipton et al., 1979], in the design of efficient graph algorithms [Lipton and Tarjan 1979] and in VLSI layout [Leighton 1983; Leiserson 1983; Valiant 1981].

Building on Lipton and Tarjan's planar separator theorem, Gilbert et al. [1984] showed that every graph with genus bounded by g has an $O(\sqrt{gn})$-separator. Another generalization was obtained by Alon et al. [1990] who showed that graphs that exclude minor isomorphic to the h-clique have an $O(h^{3/2}\sqrt{n})$-separator. Plotkin et al. [1994] reduced the dependency on h from $h^{3/2}$ to h, but in the process, they picked up a factor of $\sqrt{\log n}$. Perhaps the oldest separator result is that every tree has a 1-separator [Jordan 1869].

However, these results are apparently not applicable to geometric graphs such as nearest neighbor graphs [Preparata and Shamos 1985] when the dimension d is higher than 2.

In this paper, we give a geometrical condition of graphs embedded in d dimensions that have a small *separator*. We also present efficient linear time sequential algorithms and optimal $O(\log n)$ time parallel algorithms for finding such a small separator. The method and condition that we propose, unlike previous works, assume that the graph G comes with an embedding of its nodes in $\mathbb{R}^d$. This is a very natural assumption for nearest neighbor graphs. Our algorithm is randomized. It always splits the graph into pieces of roughly equal size, and we show that with high probability, the separator size satisfies an upper

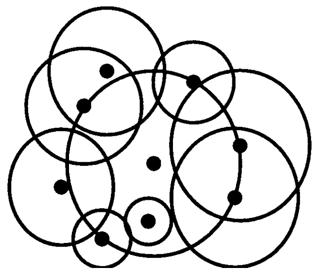

FIG. 1. A 3-ply system.

bound that is the best possible bound for this class of graphs that we consider. In conjunction with a result of Koebe [1936] that every triangulated planar graph is isomorphic to the intersection graph of a disk-packing, our result not only gives a new geometric proof of the planar separator theorem of Lipton and Tarjan [1979], but also generalizes it to higher dimensions. The separator algorithm can be used for point location and geometric divide and conquer in a fixed dimensional space.

We now review the relationship between this paper and other papers by the same authors. This paper and its companion paper [Miller et al. 1997] either extend or explain several short conference papers [Miller and Thurston 1990; Miller et al. 1991; 1993], a thesis [Teng 1991] and one journal paper [Vavasis 1991]. The focus of this paper is on graphs arising in computational geometry and the application of our separator theorem to geometric divide and conquer; the companion paper [Miller et al. 1997] focuses on finite element meshes. The authors have also jointly written a survey paper [Miller et al. 1993] that surveys the results from this paper, the companion, and several additional results by various authors on efficient centerpoint computation.

The remainder of this paper is organized as follows: In Section 2, we introduce neighborhood systems and prove our main separator theorem. We also give our randomized separator algorithm. In Section 3, we define the intersection graph of a neighborhood system and apply our main results to planar graphs and nearest neighbor graphs. In Section 4, we develop a separator-based divide-and-conquer paradigm and apply it to solve several problems in computational geometry. In Section 5, we give several open questions motivated by this research.

2. *Sphere Separators of Neighborhood Systems*

Throughout the paper we regard the dimension d as a small constant.

The class of geometric graphs that we consider is defined by the intersection of a collection of balls in d-dimensional Euclidean space. We will refer such a collection of balls as a *neighborhood system*. In this section, we will prove a geometric separator theorem for neighborhood systems. Its applications to planar graphs, k-nearest neighbor graphs, and geometric divide-and-conquer will be given in the subsequent sections.

2.1. NEIGHBORHOOD SYSTEMS

Definition 2.1.1 (*k-ply Neighborhood Systems*). A *k-ply neighborhood system in d dimensions* is a set $\{B_1, \ldots, B_n\}$ of closed balls in $\mathbb{R}^d$ such that no point in $\mathbb{R}^d$ is strictly interior to more than k of the balls.

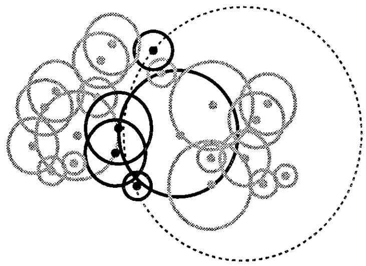

FIG. 2. The dotted circle is S. The circles with dark color are the circles in $\Gamma_O(S)$ and other circles are either in $\Gamma_E(S)$ or in $\Gamma_I(S)$.

For example, the neighborhood system given in Figure 1 is a 3-ply system.

In this definition, we used n for the number of points and d for the dimension of the embedding. We continue to use this notation throughout the paper. We also use the following notation: if $\alpha > 0$ and B is a ball of radius r, we define $\alpha \cdot B$ to be a ball with the same center as B but radius αr. In this paper, a $(d-1)$-sphere is the boundary of a d-dimensional ball.

2.2. A GEOMETRIC-SEPARATOR THEOREM. We now state our main separator theorem with respect to neighborhood systems. Each $(d-1)$-sphere S divides a neighborhood system $\Gamma = \{B_1, \ldots, B_n\}$ in $\mathbb{R}^d$ into three subsets: $\Gamma_E(S)$, the set of all balls of Γ in the exterior of S; $\Gamma_I(S)$, the set of all balls of Γ in the interior of S; and $\Gamma_O(S)$, the set of all balls of Γ that intersect S (see Figure 2).

THEOREM 2.2.1 (SPHERE-SEPARATOR THEOREM). *Suppose* $\Gamma = \{B_1, \ldots, B_n\}$ *is a k-ply system in* $\mathbb{R}^d$. *Then there is a sphere S such that*

$$|\Gamma_O(S)| = O(k^{1/d} n^{1-1/d}), \text{ and}$$

$$|\Gamma_I(S)|, |\Gamma_E(S)| \le \frac{(d+1)n}{d+2}.$$

Furthermore, for any constant ϵ *in the range* $0 < \epsilon < 1/(d+2)$ *we can compute sphere S such that* $|\Gamma_I(S)|, |\Gamma_E(S)| \le ((d+1)/(d+2) + \epsilon)n$, *and* $|\Gamma_O(S)| = O(k^{1/d} n^{1-1/d})$ *with probability at least* $1/2$.

The running time of this algorithm is bounded by $c(\epsilon, d) + O(nd)$, *where* $c(\epsilon, d)$ *is a constant depending only on* ϵ *and d.*

In this theorem and for the rest of the paper, for each finite set A, $|A|$ denotes the cardinality of A.

Remark 2.2.2. The randomization in the algorithm is over random numbers chosen by the algorithm irrespective of the input. Therefore, by rerunning the algorithm a constant number of times, we can increase the probability of success from 1/2 to $1 - \delta$ for an arbitrary $\delta > 0$.

Remark 2.2.3. The running time of the algorithm is linear, but almost all the work is done on a constant-sized subset (a subset whose size depends only on ϵ and d). The $O(nd)$ term in the running time arises from a final pass over the input to determine which of the B_i's goes into each set.

The remainder of Section 2 is devoted to proving Theorem 2.2.1.

2.3. CENTERPOINTS AND CONFORMAL MAPS. To prove Theorem 2.2.1 we need two geometric concepts: centerpoints and sphere-preserving (conformal) maps.

2.3.1. *Centerpoints*. A *centerpoint* of a given set P of points in d dimensions is a point $\mathbf{c} \in \mathbb{R}^d$ (not necessarily one of the given points) such that every hyperplane through $\mathbf{c}$ divides the given points approximately evenly (in the ratio $d : 1$ or better) [Edelsbrunner, Section 4]. It follows from Helly's theorem that every finite point set in $\mathbb{R}^d$ has a centerpoint. Various proofs can be found in Danzer et al. [1963]; Edelsbrunner [1987]; Miller et al. [1993]. It follows directly from these proofs that such a centerpoint can be found by Linear Programming on $O(n^d)$ linear inequalities of d variables. Unfortunately, no linear-time algorithm is known for computing centerpoints in higher dimensions. The following sampling algorithm can efficiently compute an approximate centerpoint:[1]

Algorithm: (*Sampling for Approximate Centerpoints*)

Input: (a point set $P \subset \mathbb{R}^d$)

1. Select a subset S of P with size l uniformly at random;
2. Compute a centerpoint $\mathbf{c}_S$ of S, using the Linear Programming algorithm for centerpoints;
3. Output $\mathbf{c}_S$.

It can be shown [Haussler and Welzl 1987; Teng 1991; Vapnik and Chervonenkis 1971] that for any constant $\epsilon < 1$, the above algorithm will compute a $(d + \epsilon) : 1$ centerpoint with high probability provided that $l > q(\epsilon, d)$, where q is a function that does not depend on n. Therefore, we can approximate a centerpoint in random constant time. In practice, we use an even simpler centerpoint approximation algorithm (see Gilbert et al. [1997], Miller et al. [1993], and Clarkson et al. [1993]).

2.3.2. *Conformal Mappings*. In our separator algorithm and our proof of Theorem 2.2.1, we map a neighborhood system from $\mathbb{R}^d$ to the unit d-sphere S^d in $\mathbb{R}^{d+1}$. An example of such a map is *stereographic projection.*

We will use notations that are consistent with our companion paper [Miller et al. 1997]. Let $\Pi(\mathbf{x})$ be the stereographic projection mapping from $\mathbb{R}^d$ to S^d. Geometrically, this map may be defined as follows: Given $\mathbf{x} \in \mathbb{R}^d$, append '0' as the final coordinate yielding $\mathbf{x}' \in \mathbb{R}^{d+1}$. Then compute the intersection of S^d with the line in $\mathbb{R}^{d+1}$ connecting $\mathbf{x}'$ to the north pole of S^d, $(0, 0, \ldots, 0, 1)^T$. This intersection point is $\Pi(\mathbf{x})$. Note that superscript T indicates transpose; thus the inner product of two vectors $\mathbf{x}$, $\mathbf{y}$ is denoted $\mathbf{x}^T\mathbf{y}$.

Algebraically, the mapping is defined as

$$\Pi(\mathbf{x}) = \begin{pmatrix} 2\mathbf{x}/\chi \\ 1 - 2/\chi \end{pmatrix} \tag{1}$$

[1] See, for example, Clarkson [1983], Haussler and Welzl [1987], Vapnik and Chervonenkis [1971], and Teng [1991].

where $\chi = \mathbf{x}^T\mathbf{x} + 1$. It is also simple to write down a formula for the inverse of Π. Let $\mathbf{u}$ be a point on S^d. Then

$$\Pi^{-1}(\mathbf{u}) = \frac{\bar{\mathbf{u}}}{1 - \mathbf{u}_{d+1}},$$

where $\bar{\mathbf{u}}$ denotes the first d entries of $\mathbf{u}$ and u_{d+1} is the last entry.

We will prove Theorem 2.2.1 using the fact that stereographic projection is *sphere-preserving*, that is, it maps spheres and hyperplanes (degenerate spheres) of $\mathbb{R}^d$ to spheres on S^d. A direct proof of this fact is given in Miller et al. [1997]. We will now give a somewhat indirect proof that stereographic projection preserves spheres. The purpose of this indirect proof is to explain some interesting properties of conformal maps in high dimensions as well as their connections back to two-dimensional conformal maps.

Define the inverse map from $\mathbb{R}^d$ to $\mathbb{R}^d$ to be

$$R(\mathbf{v}) = \frac{\mathbf{v}}{\mathbf{v}^T\mathbf{v}},$$

for all $\mathbf{v} \in \mathbb{R}^d$. We adopt the convention that $R(\vec{0}) = \infty$ and $R(\infty) = \vec{0}$, so that R is defined on $\mathbb{R}^d \cup \infty$.

Notice that each point $\mathbf{v}$ on the unit sphere S^{d-1} in $\mathbb{R}^d$ is mapped to itself. We now show that the inverse map R preserves spheres. In this statement, we will regard a hyperplane as a sphere as well.

Every sphere in $\mathbb{R}^d$ can be expressed by a quadratic equation of the following form:

$$a\mathbf{x}^T\mathbf{x} + b\mathbf{x}^T\mathbf{v}_0 + \mathbf{c} = 0,$$

where $\mathbf{v}_0$ is a constant d-vector. When $a = 0$, it is a hyperplane and when $\mathbf{c} = 0$ it is a sphere containing the origin, $\vec{0} \in \mathbb{R}^d$.

Notice that for each $\mathbf{v} \in \mathbb{R}^d$,

$$R(R(\mathbf{v})) = \frac{(\mathbf{v}/\mathbf{v}^T\mathbf{v})}{(\mathbf{v}/\mathbf{v}^T\mathbf{v})^T(\mathbf{v}/\mathbf{v}^T\mathbf{v})} = \frac{(\mathbf{v}/\mathbf{v}^T\mathbf{v})}{(1/\mathbf{v}^T\mathbf{v})} = \mathbf{v}.$$

Therefore, the inverse map is an involution.

PROPOSITION 2.3.2.1. *The inverse map preserves spheres.*

PROOF. Let $a\mathbf{y}^T\mathbf{y} + b\mathbf{y}^T\mathbf{v}_0 + c = 0$ be a sphere and let

$$\mathbf{y} = R(\mathbf{x}) = \frac{\mathbf{x}}{\mathbf{x}^T\mathbf{x}}.$$

After the substitution, we have

$$a\left(\frac{\mathbf{x}}{\mathbf{x}^T\mathbf{x}}\right)^T\left(\frac{\mathbf{x}}{\mathbf{x}^T\mathbf{x}}\right) + b\left(\frac{\mathbf{x}}{\mathbf{x}^T\mathbf{x}}\right)\mathbf{v}_0 + \mathbf{c} = 0,$$

yielding

$$a + b\mathbf{x}^T\mathbf{v}_0 + c\mathbf{x}^T\mathbf{x} = 0. \qquad \square$$

We now express stereographic projection (from $\mathbb{R}^d$ to S^d) and its inverse (from S^d to $\mathbb{R}^d$) in term as the inverse map from $\mathbb{R}^{d+1}$ to $\mathbb{R}^{d+1}$. Because S^d is embedded in $\mathbb{R}^{d+1}$, we need to first embed $\mathbb{R}^d$ in $\mathbb{R}^{d+1}$. Let L be the "natural" map from $\mathbb{R}^d$ to the hyperplane $x_{d+1} = 0$ in $\mathbb{R}^{d+1}$, that is, L sends each point $\mathbf{x} \in \mathbb{R}^d$ to a point in $\mathbb{R}^{d+1}$ by appending '0' as the final coordinate. Let $\mathbf{e} = (0, 0, \ldots, 0, 1) \in \mathbb{R}^{d+1}$. Let G be a map from $\mathbb{R}^{d+1}$ to $\mathbb{R}^{d+1}$ such that for each $\mathbf{u} \in \mathbb{R}^{d+1}$,

$$G(\mathbf{u}) = \frac{2(\mathbf{u} - \mathbf{e})}{(\mathbf{u} - \mathbf{e})^T(\mathbf{u} - \mathbf{e})} + \mathbf{e} = 2R(\mathbf{u} - \mathbf{e}) + \mathbf{e}. \tag{2}$$

Notice that G is an involution, because $G(\mathbf{u}) = 2R(\mathbf{u} - \mathbf{e}) + \mathbf{e}$, $R(G(\mathbf{u}) - \mathbf{e}) = R(2R(\mathbf{u} - \mathbf{e}))$. From the fact that $R(2\mathbf{v}) = R(\mathbf{v})/2$ for all $\mathbf{v} \in \mathbb{R}^{d+1}$, we have

$$R(G(\mathbf{u}) - \mathbf{e}) = R(2R(\mathbf{u} - \mathbf{e})) = \frac{\mathbf{R}(\mathbf{R}(\mathbf{u} - \mathbf{e}))}{2} = \frac{(\mathbf{u} - \mathbf{e})}{2},$$

implying

$$\mathbf{u} = 2R(G(\mathbf{u} - \mathbf{e})) + \mathbf{e} = G(G(\mathbf{u})).$$

Notice that for each $\mathbf{x} \in \mathbb{R}^d$, $(L(\mathbf{x}) - \mathbf{e})^T(L(\mathbf{x}) - \mathbf{e}) = \mathbf{x}^T\mathbf{x} + 1$, and thus we have

$$\Pi(\mathbf{x}) = G(L(\mathbf{x})).$$

Similarly, for each $\mathbf{y} \in S^d$,

$$\Pi^{-1}(\mathbf{y}) = L^{-1}(G(\mathbf{y})).$$

Because the inverse map and translations preserve spheres, G preserves spheres as well, implying stereographic projection Π preserves spheres. Thus, Π maps a ball (or a halfspace) of $\mathbb{R}^d$ to a cap on S^d, where a cap of S^d is the intersection of a closed halfspace in $\mathbb{R}^{d+1}$ with S^d.

2.3.3. *Conformally Mapping Centerpoints*. The translation of $\mathbb{R}^{d+1}$ by a vector $\mathbf{v}_0$ is a map that sends $\mathbf{v}$ to $\mathbf{v} - \mathbf{v}_0$. The dilation of $\mathbb{R}^{d+1}$ by a factor α is a map that sends $\mathbf{v}$ to $\alpha\mathbf{v}$. Clearly, they are all sphere-preserving. Other basic sphere preserving maps in $\mathbb{R}^{d+1}$ include the rigid rotations of $\mathbb{R}^{d+1}$, reflections, the inverse map, stereographic projection and its inverse.

LEMMA 2.3.3.1 (CENTERPOINT). *Let $P = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\}$ be a point set in $\mathbb{R}^d$. There is a sphere-preserving map Φ from $\mathbb{R}^d$ to S^d such that the origin is a centerpoint of $\Phi(P) = \{\Phi(\mathbf{p}_1), \ldots, \Phi(\mathbf{p}_n)\}$.*

PROOF. Let $\mathbf{c} \in I\!R^{d+1}$ be a centerpoint of $\Pi(P) = \{\Pi(\mathbf{p}_1), \ldots, \Pi(\mathbf{p}_n)\}$, where Π is stereographic projection. Let $U_\mathbf{c}$ be a rotation or a (Householder) reflection such that $U_\mathbf{c}(\mathbf{c}) = (0, \ldots, 0, \|\mathbf{c}\|)$, where $\|\mathbf{c}\| = \sqrt{\mathbf{c}^T\mathbf{c}}$ denotes the standard l_2 norm of $\mathbf{c}$. Clearly, $U_\mathbf{c}(\mathbf{c})$ is a centerpoint of $U_\mathbf{c} \circ \Pi(P)$.

For any positive α, let $D_\alpha = \Pi \circ (\alpha I) \circ \Pi^{-1}$, where I is the $d \times d$ identity matrix, that is, αI is a dilation map. Let $Q = U_\mathbf{c} \circ \Pi(P)$.

We now show that if $\alpha = \sqrt{(1 - \|\mathbf{c}\|)/(1 + \|\mathbf{c}\|)}$, then the center of S^d, $\vec{0}$, is a centerpoint of $D_\alpha(Q)$. Thus, $\Phi = D_\alpha \circ U_\mathbf{c} \circ \Pi$ satisfies the lemma.

We first consider the case when $d \geq 2$.

Let $\mathbf{c}' = U_\mathbf{c}(\mathbf{c}) = (0, \ldots, 0, \|\mathbf{c}\|)$. A *circle* of S^d is given by the intersection of a hyperplane in $\mathbb{R}^{d+1}$ with S^d. Let $C_{\mathbf{c}'}$ be all circles on S^d whose hyperplanes contain $\mathbf{c}'$. For each circle $H \in C_{\mathbf{c}'}$, let $D_\alpha(H)$ be the image of H under D_α. Because D_α preserves circles, $D_\alpha(H)$ is also a circle in S^d.

We first consider the circle $H_0 \in C_{\mathbf{c}'}$ whose hyperplane is normal to axis x_{d+1} (which connects the north pole with the south pole). Notice that $\Pi^{-1}(H_0)$ is a sphere in $\mathbb{R}^d$ centered at the origin. The radius of $\Pi^{-1}(H_0)$ is $\sqrt{(1 + \|\mathbf{c}\|)/(1 - \|\mathbf{c}\|)}$. By our choice of α, $D_\alpha(H_0)$ is the equator of S^d and hence is normal to axis x_{d+1} and contains the origin.

Suppose $\mathbf{q}_1$ and $\mathbf{q}_2$ are two points in $H_0 \cap S^d$. If the line segment between them is a diameter of H_0, then this line segment contains $\mathbf{c}'$. Moreover, the line segment between $D_\alpha(\mathbf{q}_1)$ and $D_\alpha(\mathbf{q}_2)$ is a diameter of $D_\alpha(H_0)$ and hence it contains the origin.

We now show for each $H \in C_{\mathbf{c}'}$ that the hyperplane of $D_\alpha(H)$ contains the origin as well. By doing this, we can conclude from definition of a centerpoint that the origin is a centerpoint of $D_\alpha \circ U_\mathbf{c} \circ \Pi(P)$.

The intersection of the hyperplanes of H and H_0 is an affine set of dimension $d - 1$. Because, we assume $d \geq 2$, this set has dimension at least 1 and contains $\mathbf{c}'$. Thus, it must contain a diameter of H_0. In other words, there exist two points $\mathbf{q}_1$ and $\mathbf{q}_2$ in $H_0 \cap H$ such that the line segment between them is a diameter of H_0. $D_\alpha(\mathbf{q}_1)$ and $D_\alpha(\mathbf{q}_2)$ are both in $D_\alpha(H_0) \cap D_\alpha(H)$, and the line segment between them contains the origin. Therefore, the hyperplane of $D_\alpha(H)$ also contains the origin.

When $d = 1$, we can embed $\mathbb{R}^1$ in $\mathbb{R}^2$. The proof as given above for $\mathbb{R}^2$ also shows that the lemma is true for $d = 1$. □

2.4. A RANDOMIZED ALGORITHM. We now present our separator algorithm. The algorithm uses randomization, and it chooses the separating sphere at random from a distribution that is carefully constructed so that the separator will satisfy the conclusions of Theorem 2.2.1 with high probability. The distribution is described in terms of sphere-preserving maps in $\mathbb{R}^{d+1}$.

Algorithm *Sphere Separator*

Input: a k-ply system $\{B_1, \ldots, B_n\}$ in $\mathbb{R}^d$ with centers $P = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\}$.

- **Project Up.** Compute $\Pi(P) = \{\Pi(\mathbf{p}_1), \ldots, \Pi(\mathbf{p}_n)\}$.
- **Find Centerpoint.** Compute a *centerpoint* $\mathbf{z}$ of $\Pi(P)$ in $\mathbb{R}^{d+1}$.
- **Conformal Map I: Rotate.** Compute an orthogonal $(d + 1) \times (d + 1)$ matrix $U_\mathbf{z}$ such that $U_\mathbf{z}(\mathbf{z}) = \mathbf{z}'$ where

$$\mathbf{z}' = (0, \ldots, 0, \|\mathbf{z}\|).$$

  Note that $\mathbf{z}'$ is a centerpoint of $U_\mathbf{z} \circ \Pi(P)$.
- **Conformal Map II: Dilate.** Let $D_\alpha = \Pi \circ (\alpha I) \circ \Pi^{-1}$, where $\alpha = \sqrt{(1 - \|z\|)/(1 + \|z\|)}$. As shown in Lemma 2.3.3.1, the origin $\vec{0}$ is a centerpoint of $D_\alpha \circ U_z \circ \Pi(P)$.
- **Find Great Circle.** Choose a random great circle C on S^d.
- **Unmap and Project Down.** Transform the great circle C to a sphere S in $\mathbb{R}^d$ by undoing the dilation, rotation, and stereographic projection:

$$S = \Pi^{-1} \circ U_\mathbf{z}^{-1} \circ D_\alpha^{-1}(C).$$

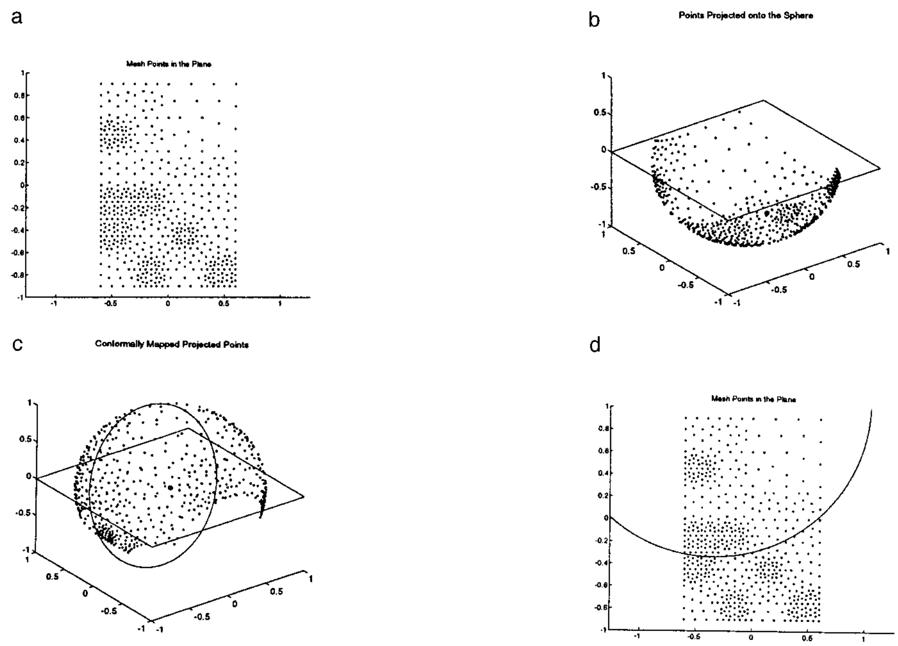

FIG. 3. (a) The point set of a neighborhood system. (b) Project Up and Find Centerpoint; the largest dot in the figure is a centerpoint. (c) Conformal Map II and Find Great Circle. (d) Unmap and Project Down. (Generated by the Matlab Geometric Separator Tool-box developed by Gilbert and Teng.

- Return S and $\Gamma_O(S)$, $\Gamma_I(S)$, and $\Gamma_E(S)$.

Figure 3 depicts the basic steps of our separator algorithm. It is generated by the Matlab Geometric Separator Tool-box of Gilbert and Teng [Gilbert et al. 1997]. The example geometric graph is generated by Eppstein. The neighborhood system of this point set is not explicitly shown in the figure. Each point defines a ball which is the largest ball centered at the point whose interior contains no other points. The ply of this neighborhood system is 6.

Notice that we can use an approximate centerpoint in the algorithm above.

In the next section, we will prove Theorem 2.2.1 by showing that $|\Gamma_O(S)| \leq O(k^{1/d}n^{1-1/d})$ with high probability.

2.5. A HIGH-LEVEL DISCUSSION. We now give a high level description of our approach to prove Theorem 2.2.1.

Let $\Gamma = \{B_1, \ldots, B_n\}$ be a k-ply neighborhood system in $\mathbb{R}^d$. Let $\mathbf{p}_i$ be the center of B_i, for $1 \leq i \leq n$. Let S^d be the unit sphere in $\mathbb{R}^{d+1}$ whose center is the origin. In the previous subsection, we have shown that there is a *sphere-preserving* map Φ from $\mathbb{R}^d$ to S^d such that the origin is a centerpoint of $\{\Phi(\mathbf{p}_1), \ldots, \Phi(\mathbf{p}_n)\}$. Let $\Phi(B_i)$ be the image of B_i on S^d; $\Phi(B_i)$ is a cap on S^d.

A *great circle* of S^d is the intersection of S^d with a hyperplane that passes through the center of S^d. Clearly, every great circle divides S^d into two open half-spheres (hemispheres). Each great circle C of S^d divides $\{\Phi(B_1), \ldots, \Phi(B_n)\}$ into three sets, C_1, C_{-1} and C_0 where C_1 and C_{-1} respectively contains those caps that are completely in each of the open hemisphere, and C_0 contains those caps that intersect C. Because the center of S^d is a centerpoint of

$\{\Phi(\mathbf{p}_1), \ldots, \Phi(\mathbf{p}_n)\}$, we have $|C_{-1}|, |C_1| \leq (d + 1)/(d + 2)n$. In order to prove Theorem 2.2.1, it is then sufficient to show that there exists a great circle C of S^d such that $|C_0| = O(k^{1/d}n^{1-1/d})$. We will prove this by arguing that the expected size of C_0 is $O(k^{1/d}n^{1-1/d})$ when C is chosen uniformly among all great circles of S^d. Since the result of this random choice is always a nonnegative number, we conclude that the probability of exceeding the expected value by more than a factor of 2 is at most 0.5.

2.6. A GEOMETRIC TECHNIQUE FOR PROVING SEPARATOR THEOREMS. A much simplified proof has been obtained, independently, by Agarwal and Pach [1995] and Spielman and Teng [1996a] since the conference publication of our result.[2] We will present our original proof here because its technique might be useful for other problems. Readers interested in the simpler proof should skip the remainder of Section 2 and refer to Pach and Agarwal [1995] and Spielman and Teng [1996a]. We will also give a high-level explanation of the simpler proof in Section 5.

Our approach is to first design a continuous function and apply a continuous version of the separator theorem (given in the next subsection) to show that some "weighted surface area" of S is "small," from which we then show the number of balls that intersect S is "small."

2.6.1. *A Continuous Separator Theorem*. Suppose $f(x)$ is a real-valued nonnegative function defined on $\mathbb{R}^d$ such that f^k is integrable for all $k = 1, 2, 3, \ldots$. Such an f is called a *cost function*. The *total volume* of the function f is defined as

$$\text{Total-Volume}(f) = \int_{\mathbf{v} \in \mathbb{R}^d} (f(\mathbf{v}))^d (\mathbf{dv})^d$$

Suppose S is a $(d - 1)$-sphere in $\mathbb{R}^d$. The *surface area* of S is then

$$\text{Area}(f, S) = \int_{\mathbf{v} \in S} (f(\mathbf{v}))^{d-1} (\mathbf{dv})^{d-1}$$

Let $P = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\}$ be a point set in $\mathbb{R}^d$. Let Φ be a sphere-preserving map from $\mathbb{R}^d$ to S^d so that the center of S^d is a centerpoint of $\Phi(P)$ (see Lemma 2.3.3.1). Recall that our sphere separator algorithm computes such a map and then uses its inverse to map a random great circle back to a sphere in $\mathbb{R}^d$.

Let S be a sphere in $\mathbb{R}^d$. The weighted surface area of S equal to Area(f, S). Because Φ carries S in $\mathbb{R}^d$ to a circle $\Phi(S)$ in S^d, we define $\text{Cost}(\Phi(S)) = \text{Area}(f, S)$. Let $\text{Avg}_\Phi(f)$ be the average cost of all great circles of S^d.

The following theorem has been proven in our companion paper [Miller et al. 1997].

THEOREM 2.6.1.1 (CONTINUOUS SEPARATOR). *Suppose f is a cost function on $\mathbb{R}^d$. Let Φ be the map from $\mathbb{R}^d$ to S^d constructed in Lemma 2.3.3.1. Then,*

$$Avg_\Phi(f) = O((\textit{Total-Volume}(f))^{1-1/d}).$$

[2] See, for example, Miller et al. [1991], Miller and Thurston [1990], and Miller and Vavasis [1991].

Remark 2.6.1.2. The conformality of the mapping from $\mathbb{R}^d$ to S^d is necessary in Theorem 2.6.1.1 in order to deal with the volume elements in the high-dimensional integrations. We refer the reader to our companion paper [Miller et al. 1997] for a discussion.

2.6.2. *Construction of a Cost Function*. To prove Theorem 2.2.1, it is sufficient to construct, for each k-ply system $\Gamma = \{B_1, \ldots, B_n\}$, a continuous function whose total volume is $O(k^{1/(d-1)}n)$ such that each sphere S in $\mathbb{R}^d$ intersects at most $O(\text{Area}(f, S))$ balls of Γ.

Let r_i be the radius of B_i and let $\gamma_i = 2r_i$, and define

$$f_i(\mathbf{x}) = \begin{cases} 1/\gamma_i & \text{if} \quad \|\mathbf{x} - \mathbf{p}_i\| \le \gamma_i \\ 0 & \text{otherwise.} \end{cases}$$

Intuitively, f_i sets up a cost on each $(d-1)$-sphere S such that the closer S is to B_i, the larger B_i contributes to the surface area of S. Using a simple geometric argument (see the companion paper [Miller et al. 1997]), we can show that for any sphere S in $\mathbb{R}^d$,

$$|\Gamma_O(S)| = O(\text{Area}_f(S)). \tag{3}$$

The function f_i is called the *local* function of B_i. We define our continuous function f as

$$f(\mathbf{x}) = l_{d-1}(f_1(\mathbf{x}), \ldots, f_n(\mathbf{x})) = \left(\sum_{i=1}^{n} (f_i(\mathbf{x}))^{d-1} \right)^{1/(d-1)},$$

where for each positive integer p, l_p denotes the standard pth norm in Euclidean space, that is, for each $a_1, \ldots, a_n$,

$$l_p(a_1, \ldots, a_n) = \left(\sum_{i=1}^{n} \left| a_i \right|^p \right)^{1/p}.$$

2.6.3. *Bounding the Total Volume of the Cost Function*. Now we give an upper bound on the total volume of f. Let V_d be the volume of a unit ball in $\mathbb{R}^d$. Clearly, $\int_{\mathbf{x} \in \mathbb{R}^d} (f_i(\mathbf{x}))^d (d\mathbf{x})^d = V_d$.

Consequently, letting

$$g(\mathbf{x}) = L_d(f_1, \ldots, f_n) = \left(\sum_{i=1}^{n} (f_i(\mathbf{x}))^d \right)^{1/d},$$

we have

$$\text{Total-Cost}(g) = \int_{\mathbf{x} \in \mathbb{R}^d} (g(\mathbf{x}))^d (d\mathbf{x})^d = V_d n$$

LEMMA 2.6.3.1. *Suppose $\Gamma = \{B_1, \ldots, B_n\}$ is a k-ply system in $\mathbb{R}^d$. Let $f_1, \ldots, f_n$ and f be functions defined as above. Then*

$$\textit{Total-Cost}(f) = O(k^{1/(d-1)}n).$$

PROOF. Because Total-Cost$(g) = V_d n$, it is sufficient to show that for all $\mathbf{x} \in \mathbb{R}^d$,

$$(f(\mathbf{x}))^d \leq c_d k^{1/(d-1)} \cdot (g(\mathbf{x}))^d.$$

We focus on a particular point $\mathbf{p} \in \mathbb{R}^d$. Notice that if $g(\mathbf{p}) = 0$, then, $f(\mathbf{p}) = 0$ as well. The inequality follows. Now, assume $g(\mathbf{p}) > 0$ and define

$$M_l = \{i \in \{1, \ldots, n\} : 2^{-l} \leq f_i(\mathbf{p}) < 2^{-l+1}\},$$

for all l such that $-\infty < l < \infty$.

Because $\cup_{-\infty \leq l \leq \infty} M_l = \{i : f_i(\mathbf{p}) \neq 0\}$ and M_l's are pairwise disjoint, each index i such that $f_i(\mathbf{p}) \neq 0$ occurs in exactly one of M_l's. Let $m_l = |M_l|$. We claim $m_l \leq 6^d k$.

We now prove the claim. For each $i \in M_l$, by the definition of M_l and f_i, $2^{l-1} \leq \gamma_i \leq 2^l$, where $\gamma_i = 2r_i$. Let B be a ball centered at $\mathbf{p}$ with radius $2^l + 2^{l-1}$. Since $\|\mathbf{p} - \mathbf{p}_i\| \leq \gamma_i$, it follows $B_i \subset B$. Because the neighborhood system has ply k, we have

$$k \cdot \mathrm{vol}(B) \geq \sum_{j \in M_l} \mathrm{vol}(B_j).$$

Let $V_d(r)$ be the volume of a ball in $\mathbb{R}_d$ of radius r. Because for all $j \in M_l$, $\mathrm{vol}(B_j) \geq V_d(2^{l-2})$,

$$k \cdot V_d(2^l + 2^{l-1}) \geq |M_l| V_d(2^{l-2}),$$

which implies $|M_l| \leq 6^d k$, completing the proof of the claim. Now, we have

$$\begin{aligned}
(f(\mathbf{p}))^d &= \left(\sum_{l=-\infty}^{\infty} \sum_{i \in M_l} f_i(\mathbf{p})^{(d-1)} \right)^{d/(d-1)} \\
&\leq \left(\sum_{l=-\infty}^{\infty} m_l (2^{-l+1})^{d-1} \right)^{d/(d-1)} \\
&\leq 2^d \left(\sum_{l=-\infty}^{\infty} m_l (2^{-l})^{d-1} \right)^{d/(d-1)},
\end{aligned}$$

where $m_l \leq 6^d k$.

We now use Inequality (4) below, established in our companion paper [Miller et al. 1997] to bound the right hand side of the equation above.

Let $\ldots, m_{-1}, m_0, m_1, m_2, \ldots$ be a doubly infinite sequence of nonnegative numbers such that each m_i is bounded above by θ and such that at most a finite number of m_i's are nonzero. Let $d \geq 2$ be an integer. Then

$$\left(\sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)} \right)^{d/(d-1)} \leq c_d \theta^{1/(d-1)} \sum_{k=-\infty}^{\infty} m_k 2^{-kd}. \tag{4}$$

where c_d is a positive number depending on d.

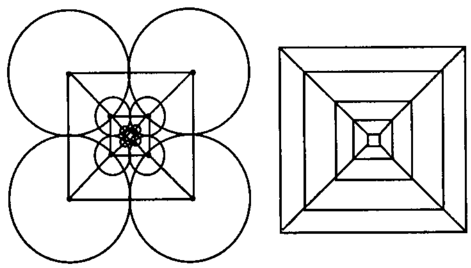

FIG. 4. Why not hyperplane separators.

Setting $\theta = 6^d k$ and applying Inequality (4), we obtain

$$f(\mathbf{p})^d \le c_d 2^d (6^d k)^{1/(d-1)} \sum_{l=-\infty}^{\infty} m_l 2^{-ld}.$$

This summation is a lower bound on $g(\mathbf{p})^d$ because for each $i \in M_l$, $f_i(\mathbf{p})^d \ge 2^{-ld}$. This concludes the proof of the lemma. □

Consequently, by Theorem 2.6.1.1, there exists a $(d + 1)/(d + 2)$-splitting sphere S of Γ with

$$\text{Area}_f(S) = O(k^{1/d} n^{1-1/d}).$$

From the definition of centerpoint, we have

$$|\Gamma_I(S)|, |\Gamma_E(S)| \le \frac{(d + 1)n}{d + 2}.$$

By Inequality (3), we have

$$|\Gamma_O(S)| = O(\text{Area}_f(S)) = O(k^{1/d} n^{1-1/d}).$$

We thus proved Theorem 2.2.1.

2.7. SPHERES VS HYPERPLANES. The simplest way to split a set of points in d-space is to use a $(d - 1)$-dimensional hyperplane. Notice that a $(d - 1)$-dimensional hyperplane is just a degenerate $(d - 1)$-sphere. Like a $(d - 1)$-dimensional sphere, a $(d - 1)$-dimensional hyperplane $\mathbf{h}$ partitions $\mathbb{R}^d$ into three subsets, $\mathbf{h}^+$, those that are above $\mathbf{h}$, $\mathbf{h}^-$, those below $\mathbf{h}$, and $\mathbf{h}$ itself, respectively. We now show, for some k-ply neighborhood systems, that it is necessary to use sphere to achieve the bound given in Theorem 2.2.1.

One such an example is given in Figure 4. It is a 2-ply system. Notice that any hyperplane that divides the neighborhood system into two of a constant ratio must intersect $\Omega(n)$ balls of the 2-ply system.

In contrast, the above 2-ply system does have a sphere that intersects $O(n^{1-1/d})$ balls and divides the rest of the balls into two sets of ratio no worst than $1 : (d + 1)$ as illustrated in Figure 5. The ratio of the two partitioned sets is $1 : 1$ in the example above.

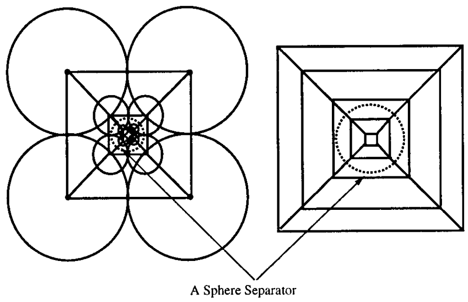

FIG. 5. Why sphere separators.

3. *Intersection Graphs*

In Section 2, we analyzed separators of an abstract geometric arrangement called a k-ply neighborhood system. The purpose of this section is to apply this abstraction to concrete classes of graphs. The most straightforward such application is to the intersection graph of k-ply systems. We also show in this section that the theory applies to planar graphs, sphere packings, and nearest neighbor graphs. Other classes of graphs not detailed here, such as finite subgraphs of the regular d-dimension grid-graph, are also covered by the theory developed in this section.

We will use the following definition of graph separator.

Definition 3.1 (*Separators*). A subset of vertices C of a graph G with n vertices is an $f(n)$-*separator* that δ-*splits* if $|C| \leq f(n)$ and the vertices of $G - C$ can be partitioned into two sets A and B such that there are no edges from A to B, and $|A|$, $|B| \leq \delta n$, where f is a function and $0 < \delta < 1$.

Given a neighborhood system, it is possible to define the intersection graph associated with the system (see Figure 6).

Definition 3.2 (*Intersection Graphs*). Let $\Gamma = \{B_1, \ldots, B_n\}$ be a neighborhood system. The *intersection graph* of Γ is the undirected graph with vertices $V = \Gamma$ and edges

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

It follows directly from Theorem 2.2.1, that the intersection graph of a k-ply neighborhood system has a small separator.

THEOREM 3.3. *Suppose $\Gamma = \{B_1, \ldots, B_n\}$ is a k-ply neighborhood system in $\mathbb{R}^d$. Then the intersection graph of Γ has an $O(k^{1/d}n^{1-1/d})$ separator that $(d+1)/(d+2)$-splits.*

The separator bound of Theorem 3.3 is the best possible in both k and n up to a constant factor. An $\Omega(n^{1-1/d})$ bound on the intersection graph of a 1-ply neighborhood system appeared in Vavasis [1991]. Let P be the set of all points of the $m \times m \times \cdots \times m$ regular grid in $\mathbb{R}^d$, where $n = m^d$. It has been shown in Teng [1991], for a sufficiently large n, that the k-nearest neighbor graph of P has no separator of size $o(k^{1/d}n^{1-1/d})$.

3.1. SPHERE-PACKINGS AND PLANAR GRAPHS. 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

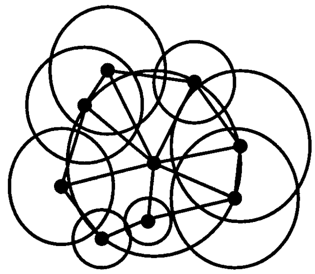

FIG. 6. The intersection graph of a 3-ply system.

a point; each edge is represented by a continuous curve connecting the two points which represent its end vertices, and no two curves share any points, except at their ends. We will only consider simple planar graphs which are graphs that do not have self-loops nor multiple edges between any pair of vertices.

We now show that Theorem 2.2.1 in conjunction with the beautiful Koebe-embedding result of planar graphs [Koebe 1936; Andreev 1970a; 1970b; Thurston 1988] gives a geometric proof of the Lipton and Tarjan planar separator theorem.

Let a *disk-packing* be a set of disks $D_1, \ldots, D_n$ that have disjoint interiors. Notice that every disk-packing is a 1-ply neighborhood system. We call the intersection graph of a disk-packing a *disk-packing graph.* It is not hard to see that every disk-packing graph is a planar graph. Koebe [1936] showed that in fact every planar graph can be represented as the intersection graph of a disk-packing. We call such a realization a *Koebe*-embedding of the planar graph. For a history of this result, including a comparison of Koebe's original result versus the Andreev–Thurston's proof, see Ziegler [1988].

THEOREM 3.1.1 (KOEBE). *Every triangulated planar graph G is isomorphic to a disk-packing graph*.

Because disk-packings are 1-ply systems in two dimensions, it follows from the existence of Koebe-embedding of planar graphs and our sphere separator Theorem 2.2.1 that every planar graph has an $O(\sqrt{n})$-separator that 3/4-splits.

We can extend the disk-packing to high dimensions: A *sphere-packing* is a neighborhood system $\Gamma = \{B_1, \ldots, B_n\}$ in $\mathbb{R}^d$ whose balls have disjoint interiors. Clearly, each sphere packing is a 1-ply neighborhood system. Therefore, the intersection graph of a sphere packing has an $O(n^{1-1/d})$ separator that $(d + 1)/(d + 2)$-splits.

Remark 3.1.2. Koebe's result strengthens Fáry's [1948] and Tutte's [1960; 1963] theorem that every planar graph can be embedded in the plane such that each edge is mapped to a straight line segment (see Thomassen [1980] and De Fraysseix et al. [1988].

Remark 3.1.3. Spielman and Teng [1996a] have recently demonstrated that the application of Theorem 2.2.1 on the Koebe-embedding of a planar graph finds a $1.84\sqrt{n}$-separator that 3/4-splits. The two constants 1.84 and 3/4 occurring in Spielman and Teng [1996a] lead to the best known bound for the constants in planar nested dissection, improving on Lipton et al. [1979] and other subsequent improved constants.

Given a Koebe-embedding, our geometric algorithm runs in random linear time with a constant smaller than that of Lipton–Tarjan's linear-algorithm [Lipton and Tarjan 1979]. It is still quite expensive to compute a Koebe embedding. Mohar [1993] has recently developed a polynomial time algorithm.

3.2. NEAREST NEIGHBOR GRAPHS. The *nearest neighbor graph* is an important class of graphs in computational geometry [Preparata and Shamos 1985]. The nearest neighbor graph arises naturally in practical applications such as image reconstruction and pattern recognition.

Let $P = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\}$ be a point set in $\mathbb{R}^d$. For each $\mathbf{p}_i \in P$, let $N_k(\mathbf{p}_i)$ be the set of k points closest to $\mathbf{p}_i$ in P (where ties are broken arbitrarily). A *k-nearest neighbor graph* [Preparata and Shamos 1985] of P is a graph with vertex set $\{\mathbf{p}_1, \ldots, \mathbf{p}_n\}$ and edge set

$$E = \{(\mathbf{p}_i, \mathbf{p}_j) : \mathbf{p}_i \in N_k(\mathbf{p}_j) \text{ or } \mathbf{p}_j \in N_k(\mathbf{p}_i)\}.$$

In this section, we show that every k-nearest neighbor graph in $\mathbb{R}^d$ is a subgraph of the intersection graph of a $O(k)$-ply neighborhood system.

For each point $\mathbf{p}_i$, let $B_i^{(k)}$ be the largest ball centered at $\mathbf{p}_i$ that contains at most k points from P, counting $\mathbf{p}_i$ itself, in the interior of $B_i^{(k)}$. Clearly the radius of $B_i^{(k)}$ is equal to the distance from $\mathbf{p}_i$ to its kth nearest neighbor in P. We call $N_k(P) = \{B_1^{(k)}, \ldots, B_n^{(k)}\}$ the k-nearest neighborhood system for P.

Notice that the k-nearest neighbor graph of P is a subgraph of the intersection graph of $\{B_1^{(k)}, \ldots, B_n^{(k)}\}$ because if $\mathbf{p}_j$ is one of $\mathbf{p}_i$'s k nearest neighbors, then $\mathbf{p}_j$ is contained in $B_i^{(k)}$ and hence $B_i^{(k)}$ must intersect with $B_j^{(k)}$.

We now show that the ply of $N_k(P)$ is at most $\tau_d k$, where τ_d is the *kissing number* in d dimensions, which is the maximum number of nonoverlapping unit balls in $\mathbb{R}^d$ that can be arranged so that they all touch a central unit ball [Conway and Sloane 1988]. It is known that $\tau_1 = 2$, $\tau_2 = 6$, $\tau_3 = 12$, $\tau_8 = 240$, and $\tau_{24} = 196560$. Although there is no explicit formula known for the kissing number τ_d for a general choice of d, it can be bounded from above and below by the following inequalities.

$$2^{0.2075..d(1+o(1))} \leq \tau_d \leq 2^{0.401d(1+o(1))}.$$

The first inequality was given by Kabatiansky and Levenshtein [1978] and the second one by Wyner [1965].

LEMMA 3.2.1 (PLY LEMMA). *Let $P = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\}$ be a point set in $\mathbb{R}^d$. Then the ply of $N_k(P)$ is bounded by $\tau_d k$.*

PROOF. Denote the balls in $N_k(P)$ by $\{B_1, \ldots, B_n\}$. We first prove the lemma when $k = 1$. In this case no ball contains the center of other balls in its interior.

Let $\mathbf{p}$ be the point in $\mathbb{R}^d$ with the largest ply. Without loss of generality, let $\{B_1, \ldots, B_t\}$ be the set of all balls that contain $\mathbf{p}$. Let C_i be the ball centered at $\mathbf{p}_i$ with radius $\|\mathbf{p}_i - \mathbf{p}\|$, and hence $\mathbf{p}$ is on the boundary of C_i for each i in the range $1 \leq i \leq t$ (see Figure 7). Clearly, C_i is contained in B_i and C_i does not contains the center of any other balls.

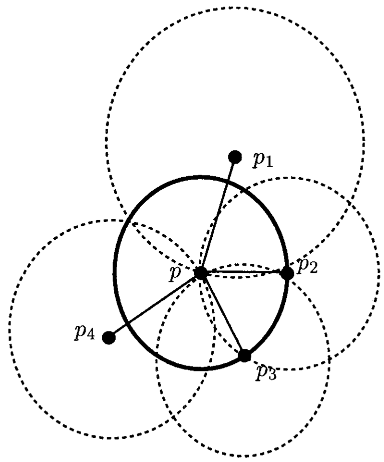

FIG. 7. The set of balls that touch a point.

Let $\delta = \min\{\|\mathbf{p} - \mathbf{p}_i\| : 0 \leq i \leq t\}$. Let $S_{\mathbf{p}}$ be the sphere centered at $\mathbf{p}$ with radius δ. Let $\mathbf{q}_i$ be the intersection of the ray $\mathbf{pp}_i$ with the sphere $S_{\mathbf{p}}$. We claim that for each pair $i, j, i \neq j$, in the range $1 \leq i, j \leq t$, $\|\mathbf{q}_i - \mathbf{q}_j\| \geq \delta$.

Without loss of generality, assume $\|\mathbf{p} - \mathbf{p}_i\| \geq \|\mathbf{p} - \mathbf{p}_j\|$. Let s be a point on the ray $\mathbf{pp}_i$ such that $\|\mathbf{p} - \mathbf{s}\| = \|\mathbf{p} - \mathbf{p}_j\|$ (see Figure 8). It follows $\|\mathbf{p} - \mathbf{p}_i\| = \|\mathbf{p} - \mathbf{s}\| + \|\mathbf{s} - \mathbf{p}_i\|$.

By the triangle inequality, we have $\|\mathbf{s} - \mathbf{p}_i\| + \|s - \mathbf{p}_j\| \geq \|\mathbf{p}_i - \mathbf{p}_j\|$.

Because $\mathbf{p}_j \notin C_i$, and the radius of C_i is $\|\mathbf{p} - \mathbf{p}_i\|$, we have $\|\mathbf{p} - \mathbf{p}_i\| \leq \|\mathbf{p}_i - \mathbf{p}_j\|$. Thus $\|\mathbf{p} - \mathbf{s}\| \leq \|\mathbf{s} - \mathbf{p}_j\|$.

By the similarity of triangles $\Delta\mathbf{pq}_i\mathbf{q}_j$ and $\Delta\mathbf{psp}_j$, we have $\|\mathbf{q}_i - \mathbf{q}_j\| \geq \|\mathbf{p} - \mathbf{q}_i\| = \delta$. Notice that the kissing number τ_d is equal to the maximum number of points that can be arranged on a unit $(d - 1)$-sphere (the boundary of a unit d-ball), such that the distance between each pair of points is at least 1. Therefore, $t \leq \tau_d$, completing the proof of the lemma when $k = 1$.

We now prove the lemma for any $k > 1$. Without loss of generality, assume $B_1, \ldots, B_t$ contain $\mathbf{p}$. Define a subset Q of $\{\mathbf{p}_1, \ldots, \mathbf{p}_t\}$ by the following procedure. Initially, let $P = \{\mathbf{p}_1, \ldots, \mathbf{p}_t\}$ and $Q = \emptyset$.

while $P \neq \emptyset$

(1) Let $\mathbf{q}$ be the point in P with the largest $\|\mathbf{q} - \mathbf{p}\|$, let $Q = Q \cup \{\mathbf{q}\}$;

(2) Let $P = P - \text{int}(B_{\mathbf{q}})$, (where $B_{\mathbf{q}}$ stands for the closed ball centered at $\mathbf{q}$).

Because no ball contains more than k points from $\{\mathbf{p}_1, \ldots, \mathbf{p}_t\}$ in its interior, we have $m \geq \lceil t/k \rceil$, where m denotes $|Q|$.

We now show that for all $\mathbf{q} \in Q$, $\text{int}(B_q) \cap Q = \{\mathbf{q}\}$. Suppose $Q = \{\mathbf{q}_1, \ldots, \mathbf{q}_m\}$ such that for all $i < j$, $\mathbf{q}_i$ is put Q in the above procedure before $\mathbf{q}_j$. Notice that for all $j > i$, $\mathbf{q}_j \notin \text{int}(B_{\mathbf{q}_i})$. Also because $\|\mathbf{q}_i - \mathbf{q}_j\| \geq \|\mathbf{q}_i - \mathbf{p}\| \geq \|\mathbf{q}_j - \mathbf{p}\|$, we have for all $i < j$, $\mathbf{q}_i \notin \text{int}(B_{\mathbf{q}_j})$. So $\text{int}(B_{\mathbf{q}}) \cap Q = \{\mathbf{q}\}$. Thus, $m \leq \tau_d$ which implies $t \leq \tau_d k$. □

Consequently, by Theorem 2.2.1:

THEOREM 3.2.2. *Every k-nearest neighbor graph in d dimensions has an* $O(k^{1/d}n^{1-1/d})$ *separator that* $(d + 1)/(d + 2)$*-splits.*

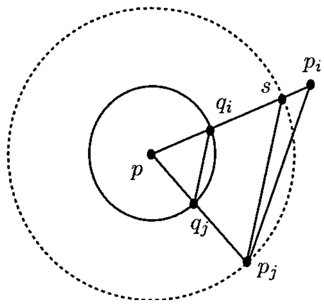

FIG. 8. The distance between q_i and q_j.

The following is an interesting consequence of the Ply Lemma 3.2.1:

COROLLARY 3.2.3. *The degree of all k-nearest neighbor graphs in d dimensions is bounded above by* $(\tau_d + 1)k$.

PROOF. For each point $\mathbf{p}_i$ in a given set P, let B_i be the largest ball centered at $\mathbf{p}_i$ such that the interior of B_i contains no more than k points from P. Notice that if $(\mathbf{p}_i, \mathbf{p}_j)$ is an edge in a k-nearest neighbor graph of P, then either $\mathbf{p}_i \in B_j$ or $\mathbf{p}_j \in B_i$. By Lemma 3.2.1, we have the degree of k-nearest neighbor graphs is bounded above by $(\tau_d + 1)k$. □

Remark 3.2.4. Toussaint [1988] called the intersection graph of $N_k(P)$ the *kth sphere-of-influence graph* of P and showed that this class of graphs can be used in image processing. It follows from Lemma 3.2.1 and Theorem 3.3.1 that the kth sphere-of-influence graph of any set of n points in $\mathbb{R}^d$ has an $O(k^{1/d}n^{1-1/d})$ separator that $(d + 1)/(d + 2)$-splits.

3.3. INDUCTIVITY OF INTERSECTION GRAPHS. In this subsection, we show that the intersection graph of a k-ply neighborhood system has a linear number of edges. Because the star graph can be realized as the intersection graph of some 1-ply neighborhood system, the maximum degree of these intersection graphs can be unbounded.

For any integer δ, a graph is δ-*inductive* if its vertices can be numbered such that each vertex has at most δ edges to higher numbered vertices. Clearly, a δ-inductive graph with n vertices has at most $(\delta \cdot n)$ edges. For example, every tree is 1-inductive and each simple planar graph is 5-inductive. The latter can be shown by observing that each planar graph has at least one vertex of degree less than 6 (by Euler's formula). So a 5-inductive numbering can be obtained by assigning the smallest number to such a vertex and inductively numbering other vertices.

THEOREM 3.3.1. *The intersection graph of a k-ply neighborhood system in* $\mathbb{R}^d$ *is* $3^d k$*-inductive.*

To prove the theorem, it is sufficient to show that each ball in a k-ply neighborhood system in $\mathbb{R}^d$ intersects at most $3^d k$ other balls of larger or equal radius. Then, we can number the balls by sorting the radius of the balls in the increasing order. Therefore, Theorem 3.3.1 follows directly from the following lemma.

LEMMA 3.3.2 (BALL INTERSECTION). *Suppose* $\Gamma = \{B_1, \ldots, B_n\}$ *is a k-ply system in* $\mathbb{R}^d$. *Then for each d-dimensional ball B with radius r,* $|\{i : B_i \cap B \neq \emptyset \text{ and } r_i \geq r\}| \leq 3^d k$.

PROOF. Without loss of generality, let $B_1, \ldots, B_t$ be the set of all balls in Γ of radius at least r that intersect B. For each i in the range $1 \leq i \leq t$, if $\mathbf{p}_i$, the center of B_i, is in $2 \cdot B$, let B'_i be the ball of radius r centered at $\mathbf{p}_i$; if $\mathbf{p}_i$ is not in $2 \cdot B$, let $\mathbf{p}'_i$ be the point common of the ray $\mathbf{pp}_i$ and the boundary of $2 \cdot B$, and let B'_i be the ball centered at $\mathbf{p}'_i$ and of radius r. In either case, $B'_i \subseteq B_i$ and B'_i intersects B, and if B_i is replaced by B'_i, the ply of the resulting neighborhood system does not increase and thus is bounded above by k. Notice that each ball $B'_i(1 \leq i \leq t)$, is contained in the ball $3 \cdot B$. We have

$$\sum_{i=1}^{t} \text{Volume}(B'_i) \leq k\text{Volume}(3 \cdot B),$$

which implies $t \leq 3^d k$. $\square$

Remark 3.3.3. Lemma 3.3.2 was given in the conference publication of this work [Miller et al. 1991] and a proof appeared in Teng's dissertation [1991]. A few years later, a variant of it was independently proved by Eppstein and Erickson [1994].

Notice that any k-inductive graph G is $(k + 1)$-colorable by the following greedy algorithm. Suppose the vertex set of G labeled by a k-inductive labeling $\{1, \ldots, n\}$. Color the vertices $n - k, \ldots, n$ by colors $1, \ldots, k + 1$, respectively. We color the remainder of the vertex set in the order of $n - k - 1, \ldots, 1$. Because each vertex is connected to at most k vertices of higher labels, we can always assign it a color that is not used by its neighbors with higher labels. So this greedy algorithm is guaranteed to use $k + 1$ colors.

COROLLARY 3.3.4. *The intersection graph of a k-ply neighborhood system in* $\mathbb{R}^d$ *is* $(3^d k + 1)$-*colorable.*

4. *Geometric Divide and Conquer*

In this section, we present a divide-and-conquer paradigm that uses Theorem 2.2.1. We will demonstrate the usefulness of this paradigm in computational geometry. The new paradigm is compared with a commonly used paradigm for solving geometry problems, the multi-dimensional divide and conquer of Bentley [1980]. We will show that this paradigm outperforms multi-dimensional divide and conquer on various geometry problems. The new paradigm also provides a good support for designing efficient parallel algorithms for geometry problems in fixed dimensions (see Frieze et al. [1992]).

4.1. POINT LOCATION. The point location problem for a neighborhood system can be defined as: given a neighborhood system $\Gamma = \{B_1, \ldots, B_n\}$ in d-space, preprocess the input to organize it into a search structure so that queries of the form "output all neighborhoods that contain a given point $\mathbf{p}$" can be answered efficiently.

Like other geometry query problems, there are three costs associated with this point location problem: the *preprocessing* time $T(n, d)$ required to build the search structure, the query time $Q(n, d)$ required to answer a query, and the *space* $S(n, d)$ required to represent the search structure in memory.

If Γ is an arbitrary neighborhood system, then there may exist some point $\mathbf{p}$ that is covered by $\Omega(n)$ balls. In this case, just to print the output would require $\Omega(n)$ work. However, if Γ is restricted to be a k-ply system, then the number of balls in the output is bounded by k. Using separator based divide and conquer, we are able to construct a search structure with the following properties:

$$T(n, d) = \text{random } O(n \log n),$$

$$Q(n, d) = O(k + \log n),$$

$$S(n, d) = O(n).$$

By saying an algorithm runs in random $t(n)$ time, we mean that the algorithm never gives a wrong output but may not terminate in the claimed time bound. The probability of success, namely, that it produces a correct output in $t(n)$ steps, is at least $1 - \lambda$ for any $\lambda > 0$. To simplify the discussion, in the following sections, we assume that the ply k is a constant.

The main idea is to use a sphere separator which intersects an $O(k^{1-\beta}n^{\beta})$ number of balls for any constant ($\beta < 1$) to partition the neighborhoods into two subsets of roughly equal size, and then recursively build search structures for each subsets.

Given a neighborhood system Γ with ply k, we will build a binary tree of height $O(\log n)$ to guide the search in answering a query. Associated with each leaf of the tree is a subset of neighborhoods in Γ, and the search structure has the property that for all $\mathbf{p} \in \mathbb{R}^d$, the set of neighborhoods that covers $\mathbf{p}$ can be found in one of the leaves.

In the following construction, we will use sphere separators that have the following useful properties

– It can be represented with $O(1)$ space.
– It takes $O(1)$ time to test whether a point is in the interior or the exterior of the sphere.

The algorithm is very simple. It first finds a sphere S that intersects $c \cdot k^{1-\beta}n^{\beta}$ balls that δ-split Γ. In the remainder of this section we assume that β, δ and c are constants with the property that $0 < \beta < 1$, $0 < \delta < 1$ and c is a positive real that only depends on the dimension d, β and δ. To apply Theorem 2.2.1, we can use $\beta = 1 - 1/d$, $\delta = (d + 1)/(d + 2) + \epsilon$ for any constant ϵ in the range $0 < \epsilon < 1/(d + 2)$, and c is the constant term of the separator size given in Theorem 2.2.1. Because testing whether a sphere intersects more than $c \cdot k^{1-\beta}n^{\beta}$ balls can be done in linear time, our randomized separator algorithm can guarantee the quality of its separators.

Let Γ_0 be the subset of balls which intersect either S or the interior of S, and Γ_1 the subset of balls which intersect either S or the exterior of S. Clearly $|\Gamma_0|$, $|\Gamma_1| \le \delta n + c \cdot k^{1-\beta}n^{\beta}$, and $|\Gamma_1| + |\Gamma_2| \le n + c \cdot k^{1-\beta}n^{\beta}$. We store the information about S, namely its center and radius, in the root of the search tree, and recursively build binary search trees for Γ_0 and Γ_1, respectively. The roots of the tree for Γ_0 and Γ_1 are respectively the left and right children of the node associated with S. The recursive construction stops when the subset has cardinality smaller than $m_0 = \alpha k$ for a constant α that depends on β, δ, c but not n and k. The precise requirements for α (and hence m_0) will be determined below.

To answer a query when given a point $\mathbf{p} \in \mathbb{R}^d$, we first check $\mathbf{p}$ against S, the sphere separator associated with the root of the search tree. There are three cases:

Case 1. If $\mathbf{p}$ is in the interior of S, then recursively search on the left subtree of S;
Case 2. If $\mathbf{p}$ is in the exterior of S, then recursively search the right subtree of S;
Case 3. If $\mathbf{p}$ is on S, then recursively search on the left subtree of S.

When reaching a leaf, we then check $\mathbf{p}$ against all balls associated with the leaf and print all those that cover $\mathbf{p}$.

The correctness of the search structure and the above searching procedure is obvious and can be proved by induction: if $\mathbf{p}$ is in the interior (exterior) of S, then all balls that cover $\mathbf{p}$ must intersect either S or the interior (exterior) of S, and hence are in the left (right) subtree of S. The time complexity to answer a query is bounded by $O(h_k(n) + m_0)$, where $h_k(n)$ is the worse-case height of the search tree for n k-ply balls. We can bound $h_k(n)$ from above by the following recurrence.

$$h_k(m) \leq \begin{cases} 1 & \text{if} \quad m \leq m_0 \\ h_k(\delta m + ck^{1-\beta}m^{\beta}) + 1 & \text{if} \quad m \geq m_0. \end{cases} \tag{5}$$

The following lemma gives an upper bound on $h_k(n)$. In proof, we will give the first condition on $m_0 = \alpha k$.

LEMMA 4.1.1. *Let h_k be a function defined above. Then $h_k(n) = O(\log n)$ for a sufficiently large constant α that depends only on d, δ, β, and c.*

PROOF. We will choose $m_0 = \alpha k$ such that for all $m \geq m_0$, $ck^{1-\beta}m^{\beta} \leq ((1 - \delta)/2)m$. This condition is true if

$$\alpha \geq \left(\frac{2c}{1 - \delta}\right)^{1/(1-\beta)}.$$

Because $h_k(m)$ is a nondecreasing function in m, we have

$$h_k(m) \leq \begin{cases} 1 & \text{if} \quad m \leq m_0 \\ h_k\left(\left(\frac{1 + \delta}{2}\right)m\right) + 1 & \text{if} \quad m > m_0. \end{cases}$$

Since $\delta < 1$ and hence $2/(1 + \delta) > 1$, we can infer $h_k(n) = \lceil \log_{2/(1+\delta)} n \rceil = O(\log n)$. □

Consequently,

$$Q(n, d) = O(\log n + m_0) = O(\log n + k).$$

We now analyze the space requirement of the search structure. First, observe that each internal node requires a constant amount of space and each leaf requires $O(m_0)$ space. To bound the total space, it is sufficient to bound the total number of leaves in the tree.

Let $s_k(m)$ denote the maximum number of leaves in the search tree for m balls. The sphere separator decomposes the data structure for m balls into two substructures, one for those balls intersecting the interior of the sphere and one for those balls intersecting the exterior of the sphere. Our separator results guarantee that number of balls in each two substructures is no more than $\delta m + ck^{1-\beta}m^{\beta}$ and the sum of number of balls from both sides is no more than $m + ck^{1-\beta}m^{\beta}/2$. Hence, there is a δ_1 such that (1) $1 - \delta_1 \le \delta_1 \le \delta$; and (2) the smaller side (either interior or exterior) has no more than $(1 - \delta_1)m$ balls. The number of balls in the larger side is at most $\delta_1 m + ck^{1-\beta}m^{\beta}$. Notice that we implicitly charge the additional term for the separator (which is bounded by $ck^{1-\beta}m^{\beta}$ in $(1 - \delta_1)m$). Thus, $s_k(m)$ is given by the following recurrence.

$$s_k(m) \le \begin{cases} 1 & \text{if} \quad m \le m_0 \\ s_k(\delta_1 m + ck^{1-\beta}m^{\beta}) + s_k((1-\delta_1)m) & \text{if} \quad m > m_0. \end{cases} \tag{6}$$

The following lemma gives an upper bound on $s_k(n)$. In proof, we will give the second condition on $m_0 = \alpha k$.

LEMMA 4.1.2. *Let s_k be the function defined above. Then $s_k(n) = O(n/k)$ for a sufficiently large constant α that depends only on d, δ, β, and c.*

PROOF. For any constant γ such that $\beta < \gamma < 1$, we use induction to establish $s_k(n) \le C(n/k - (n/k)^{\gamma})$ for a sufficiently large constant α that depends only on d, δ, β, and c and an appropriate choice of the constant $C > 1$.

Because $s_k(m) = 1$ for $m \le m_0 = \alpha k$, we need to choose C, and m_0 such that $C(m_0/k - (m_0/k)^{\gamma}) \ge 1$. Because $\gamma < 1$, this condition holds for a sufficiently large α, establishing the base for the induction. Now assuming the lemma is true for all $m < n$, using the substitution method of Cormen et al. [1990], we have

$$\begin{aligned} s_k(n) &\le s_k(\delta_1 n + ck^{1-\beta}n^{\beta}) + s_k((1-\delta_1)n) \\ &\le \frac{C(\delta_1 n + ck^{1-\beta}n^{\beta})}{k} + C\frac{((1-\delta_1)n)}{k} - C\left(\frac{(\delta_1 n + ck^{1-\beta}n^{\beta})}{k}\right)^{\gamma} - C\left(\frac{(1-\delta_1)n}{k}\right)^{\gamma} \\ &= \frac{Cn}{k} - C\left(\frac{n}{k}\right)^{\gamma} + C\left(\frac{n}{k}\right)^{\gamma} + Cc\left(\frac{n}{k}\right)^{\beta} - C\left(\delta_1\frac{n}{k} + c\left(\frac{n}{k}\right)^{\beta}\right)^{\gamma} - C\left(\frac{(1-\delta_1)n}{k}\right)^{\gamma} \\ &\le \frac{Cn}{k} - C\left(\frac{n}{k}\right)^{\gamma} + C\left(\frac{n}{k}\right)^{\gamma} + Cc\left(\frac{n}{k}\right)^{\beta} - C\left(\frac{\delta_1 n}{k}\right)^{\gamma} - C\left((1-\delta_1)\frac{n}{k}\right)^{\gamma} \\ &\le C\left(\frac{n}{k} - \left(\frac{n}{k}\right)^{\gamma}\right), \end{aligned}$$

as long as we choose C and $m_0 = \alpha k$ such that for all $n \geq m_0$,

$$C\left(\frac{\delta_1 n}{k}\right)^{\gamma} + C\left((1-\delta_1)\frac{n}{k}\right)^{\gamma} - C\left(\frac{n}{k}\right)^{\gamma} - Cc\left(\frac{n}{k}\right)^{\beta} \geq 0\,. \tag{7}$$

By a Taylor expansion of $(1 - x)^{\gamma}$ around the point 0, we have

$$(\delta_1)^{\gamma} + (1-\delta_1)^{\gamma} \geq (\delta_1)^{\gamma} + 1 - \gamma\delta_1.$$

Because $0 < \gamma < 1$ and $1 < \delta_1 < \delta < 1$, the above inequality (7) holds if

$$C[(\delta^{\gamma} - \gamma\delta)(n/k)^{\gamma} - c(n/k)^{\beta}] \geq 0. \tag{8}$$

Because $0 < \beta < \gamma < 1$, inequality (8) holds for all

$$n \geq m_0 \geq \left(\frac{C}{\delta^{\gamma} - \gamma\delta}\right)^{1/(\gamma-\beta)} k.$$

Therefore, the lemma is true for sufficiently large constants α and C that only depend on d, δ, β, and c. □

Because the number of nodes in a proper binary tree is no more than twice the number of leaves, we have for each sufficiently large constant α satisfying the conditions given in the proof of Lemma 4.1.2,

$$s(n) = O\left(\frac{n}{k}\right).$$

Therefore, the total space requirement of the above search structure is bounded by

$$S(n, d) = O(ks_k(n)) = O(n).$$

Now let us look at the time required in building such a search structure.

From Theorem 2.2.1, each k-ply system of m balls in $\mathbb{R}^d$ has a sphere separator that intersects $O(k^{1/d}n^{1-1/d})$ balls and $(d + 1)/(d + 2)$-splits the system. If we could compute such a sphere separator in deterministic $O(n)$ time, then the worst-case time required in computing such a search structure, $T_k(m)$, would be given by the following recurrence.

$$T(m) \leq \begin{cases} 1 & \text{if} \quad m \leq m_0 \\ T_k(\delta_1 m + ck^{1-\beta}m^{\beta}) + T_k((1-\delta_1)m) + O(m) & \text{if} \quad m > m_0, \end{cases} \tag{9}$$

where $\delta_1 \leq \delta$.

The following lemma gives an upper bound on $T_k(n)$. In proof, we will give the third condition on $m_0 = \alpha k$:

LEMMA 4.1.3. *Let T_k be the function defined above. Then $T_k(n) = O(n \log n)$ for a sufficiently large constant α that depends only on d, δ, β, and c.*

PROOF. We use induction to establish $T_k(n) \leq Cn \log n$ for an appropriate choice of the constant $C > 1$.

Clearly, $T_k(m) = 1 \le C$ for $m \le m_0$. This is the base of the induction. Now, assuming the lemma is true for all $m < n$, we have

$$
\begin{aligned}
T_k(n) &\le T_k(\delta_1 n + ck^{1-\beta}n^\beta) + T_k((1-\delta_1)n) + c_2 n \\
&\le C(\delta_1 n + ck^{1-\beta}n^\beta)\log(\delta_1 n + ck^{1-\beta}n^\beta) + C((1-\delta_1)n)\log((1-\delta_1)n) + c_2 n \\
&\le Cn \log(\delta n + ck^{1-\beta}n^\beta) + Cck^{1-\beta}n^\beta \log(\delta + ck^{1-\beta}n^\beta) + c_2 n \\
&\le Cn \log\left(\frac{1+\delta}{2} n\right) + Cck^{1-\beta}n^\beta \log(\delta n + ck^{1-\beta}n^\beta) + c_2 n \\
&= Cn \log n - Cn \log\left(\frac{2}{1+\delta}\right) + Cck^{1-\beta}n^\beta \log(\delta + ck^{1-\beta}n^\beta) + c_2 n \\
&\le Cn \log n,
\end{aligned}
$$

as long as we choose C and $m_0 = \alpha\, k$ such that for all $n \ge m_0$,

$$Cn \log\left(\frac{2}{1+\delta}\right) - Cck^{1-\beta}n^\beta \log(\delta + ck^{1-\beta}n^\beta) - c_2 n \ge 0. \tag{10}$$

Because $\beta < 1$ and $\delta < 1$, inequality (10) holds for sufficiently large α and $C > c_2$ which only depend on d, δ, β, and c. □

Consequently,

$$T_k(n, d) = O(n \log n).$$

Notice that, however, our algorithm is randomized. As shown in our main theorem, if $\beta = 1 - 1/d + \epsilon$ for some constant ϵ such that $0 < \epsilon < 1/d$, then the probability such a randomized algorithm outputs a sphere separator that intersects $O(k^{1-\beta}n^\beta)$ balls is at least $1 - (1/n^\epsilon)$. Moreover, in linear time, we can check whether the number of balls a sphere separator intersects is $O(k^{1-\beta}n^\beta)$. Frieze et al. [1992] in a parallel extension of this algorithm, shown that for a sufficiently large constant m_0 the search structure can be constructed in random $O(n \log n)$ time with a probability of success $1 - 1/n$.

4.2. CONSTRUCTING INTERSECTION GRAPHS. The problem of this section is to construct the intersection graph of a given neighborhood system. There is a simple solution for this problem: test each pair of balls to decide whether they intersect. Since there are $O(n^2)$ pairs and the testing of each pair can be performed in constant time, the whole construction can be performed in $O(n^2)$ time. If we require the algorithm to report all edges of the intersection graph, the above algorithm is optimal if we are working with general neighborhood systems. This is because that there could be as much as $\Omega(n^2)$ number of edges in some intersection graphs. However, every k-intersection graph has at most $O(kn)$ edges. We present a randomized $O(kn + n \log n)$ time construction algorithm.

To illustrate the idea, let us view the graph construction problem as a search problem: a problem of exploring the structure of an unknown graph with the help of some oracles. First, assume that n, the number of vertices, is known in advance and we have an oracle – the *edge oracle* – which answers the question of the form "is there an edge between vertex u and v?" in constant time. It is not hard to see that even though the number of edges is known in advance, $\Omega(n^2)$ queries have to be asked in the worst case.

Now suppose that there is more information available: it is known in advance that the graph has a $ck^{1-\beta}n^{\beta}$-separator that δ-splits and moreover each subgraph of $m > m_0$ vertices also has a $ck^{1-\beta}m^{\beta}$-separator that δ-splits, for some constant c, α, and $0 < \beta < 1$, where $m_0 = \alpha k$. Can the number of queries be reduced? It is remain to be seen whether this is true.

Now, suppose in addition, we have an oracle – a *separator oracle* – which, when presented with a subset of m vertices, delivers three sets, A, B, and C, where C is an $ck^{1-\beta}m^{\beta}$-separator that δ-splits the subgraph induced by those m vertices into A and B. Then it is sufficient to consult with the oracle $O(n)$ times to compute the structure of the unknown graph G, if k is much smaller than n.

The strategy is divide and conquer. We first present the separator oracle with the whole set of vertices and get back from the oracle three sets A, B, C, where C is a $ck^{1-\beta}n^{\beta}$-separator that δ-splits G into A and B. We then recursively search the structure of subgraphs induced by $A \cup C$ and $B \cup C$ until the size of subproblems is below m_0. Finally, we use the edge oracle to complete the graph.

The total number of query $q_k(n)$ to the separator oracle is clearly given by the following recurrence.

$$q_k(n) \leq \begin{cases} 0 & \text{if} \quad n \leq m_0 \\ q_k(\delta_1 n + ck^{1-\beta}n^{\beta}) + q((1-\delta_1)n) + 1 & \text{if} \quad n > m_0, \end{cases}$$

where $\delta_1 \leq \delta$.

By a similar argument as Lemma 4.1.2, it can be shown $q_k(n) = O(n/k)$.

Now suppose each query to the separator oracle costs $O(m)$ time, where m is the size of query. It is not hard to see that the total time $T_k(n)$ needed to search the structure of the graph is given by the following recurrence.

$$T_k(n) \leq \begin{cases} O(1) & \text{if} \quad n \leq m_0 \\ T_k(\delta_1 n + ck^{1-\beta}n^{\beta}) + T_k((1-\delta_1)n) + O(n) & \text{if} \quad n > m_0. \end{cases}$$

By Lemma 4.1.3, we have $T_k(n) = O(n \log n)$.

The divide-and-conquer algorithm for constructing intersection graphs is based on the following interesting observation: We do not need to have the intersection graph in order to compute a small separator efficiently – all we need is the neighborhood system!

To see this, let us recall how we compute a small separator of an intersection graph. First, we find a sphere separator S of low cost. This step involves computing an approximate centerpoint and a conformal map. We then compute a vertex separator from the sphere separator S. The rule of choosing vertices is very simple: If ball B_i has a common point with S, then the vertex corresponding to B_i is placed in the separator. The time complexity of the above step is $O(n)$.

THEOREM 4.2.1. *The intersection graph of a k-ply system in* $\mathbb{R}^d$ *can be computed in random O(kn + n log n) time. Moreover, the algorithm uses O(n)-space.*

In contrast, one can only derive an $O(kn \log^{d-1} n)$ time, $O(n \log^{d-1} n)$-space algorithm for computing an intersection graph using the multidimensional divide-and-conquer paradigm.

Remark 4.2.2. Guibas et al. [1994] gave an

$$O(n^{2-\frac{2}{(1+\lfloor (d+2)/2 \rfloor)+o(1)}} + kn \log^2 n)$$

time deterministic algorithm for constructing the kth sphere of influence graph of a point set in d dimensions. Our construction, using randomization, reduces the time complexity to expected $O(kn + n \log n)$.

Remark 4.2.3. Eppstein et al. [1993] gave a deterministic linear time algorithm for finding a small cost sphere separator of a k-ply system. Therefore, in theory, all results presented in this section can be made deterministic. However, the randomized construction presented in this section is much faster in practice. More practical deterministic algorithms are desirable.

Eppstein et al. also showed how to apply this divide-and-conquer method to approximate the ply k. Hence, we can apply separator based divide and conquer to neighborhood systems without knowing its ply k a-priori.

5. *Final Remarks and Open Questions*

Recently, using the duality on S^d suggested in Teng [1991] and Eppstein et al. [1993], Agarwal and Pach [1995] and Spielman and Teng [1996a], independently, gave a much simpler geometric proof that bounds the expected number of balls of a k-ply systems on S^d that a random great circle intersects (a la Sections 2.6.2 and 2.6.3): Let $\Gamma = (B_1, \ldots, B_n)$ be a k-ply neighborhood system in R^d, and let Φ be the sphere-preserving map used in Lemma 2.3.3.1. let r_i be the radius of $\Phi(B_i)$. As shown in Teng [1991], Eppstein et al. [1993], Pach and Agarwal [1995], and Spielman and Teng [1996a], it follows from the duality between points and great circles of S^d that the expected number of caps of $\Phi(\Gamma)$ that a random great circle of S^d intersect is equal to

$$c_1 \sum_{i=1}^{n} r_i,$$

for a constant c_2 depends only on d. Because $\Phi(\Gamma)$ is k-ply, the total volume of $\{\Phi(B_1), \ldots, \Phi(B_n)\}$ is at most kA_d, where A_d is the surface area of a unit d-sphere in $\mathbb{R}^{d+1}$. Therefore, there is a constant c_2 depending only on d such that

$$\sum_{i=1}^{n} r_i^d \leq c_2 k.$$

By a convexity argument or Lagrange's method, one can show that

$$c_1 \sum_{i=1}^{n} r_i = O(k^{1/d} n^{1-1/d}).$$

An important problem is, given a graph without an embedding, can its nodes be embedded in $\mathbb{R}^d$ to make it a subgraph of an intersection graph of a k-ply system?

Recently, Linial et al. [1995] studied the problem of embedding graphs in the Euclidean space so that (1) the dimension is kept as small as possible, and (2) the distances among vertices of the graph are closely matched with the distances between their geometric images. They showed that such an embedding of a graph in $\mathbb{R}^d$ in conjunction with the partitioning technique of this paper implies that the graph has a separator of size $O(n^{1-1/d})$.

In contrast to our approach, the results of Linial et al. [1995] can be applied to any graph, not only graphs arising from neighborhood systems. On the other hand, their algorithm in general requires the embedding dimension to be as large as $\Omega(\log n)$, so the bounds they attain ours are weaker than ours. It would be very interesting if there were an embedding algorithm for general graphs that would be able to find very low-dimensional embedding in the special case that the graph admits such an embedding.

It is interesting to point out that Vavasis [1991] has defined a class of geometric graphs called local graphs and showed that any local graph of n vertices has a hyperplane based separator of size $n^{1-1/d}$. The class of local graphs is properly contained in the class of overlap graphs defined in our companion paper [Miller et al. 1996] and is much weaker than k-intersection graphs. Previously, hyperplanes have been used in the recursive coordinate bisection heuristic.

Recently, Plotkin et al. [1994] improved the bounded forbidden minor separator theorem of Alon et al. [1990]. Their work, in conjunction with a structure lemma of Teng [1994], gives a combinatorial proof of a much weaker version of the separator theorem for the class of intersection graphs presented in this paper.

It is also interesting to determine how our algorithm performs in practice compared to other current algorithms such as the spectral method. This is the subject of current work by Gilbert et al. [1997]. Very recently, using Koebe-embedding and sphere-preserving mapping in a similar way to the work of this paper, Spielman and Teng [1996b] showed that the spectral method can be used to find edge bisectors of size $O(\sqrt{n})$ for k-nearest neighbor graphs and bounded degree planar graphs.

ACKNOWLEDGMENTS. We would like to thank David Applegate, Marshall Bern, David Eppstein, Alan Frieze, John Gilbert, Bruce Hendrickson, Ravi Kannan, Tom Leighton, Mike Luby, Oded Schramm, Dan Spielman, Dafna Talmor, Doug Tygar, and Kim Wagner for invaluable help and discussions. We would like to thank both referees for their very helpful comments and suggestions.

REFERENCES

ALON, N., SEYMOUR, P., AND THOMAS, R. 1990. A separator theorem for graphs with an excluded minor and its applications. In *Proceedings of the 22th Annual ACM Symposium on Theory of Computing* (Baltimore, Md., May 14–16). ACM, New York, pp. 293–299.

ANDREEV, E. M. 1970a. On convex polyhedra in Lobacevskii space. *Math. USSR Sbornik 10*, 3, 413–440.

ANDREEV, E. M. 1970b. On convex polyhedra of finite volume in Lobacevskii space. *Math. USSR Sbornik, 12*, 2, 270–259.

BENTLEY, J. L. 1980. Multidimensional divide-and-conquer. *Commun. ACM 23*, 4 (Apr.), 214–229.

CLARKSON, K. 1983. Fast algorithm for the all-nearest-neighbors problem. In *Proceedings of the 24th Annual Symposium on Foundations of Computer Science*. IEEE, New York, pp. 226–232.

CLARKSON, K., EPPSTEIN, D., MILLER, G., STURTIVANT, C., AND TENG, S.-H. 1993. Approximating center points with iterated radon points. In *Proceedings of 9th Annual ACM Symposium on Computational Geometry* (San Diego, Calif., May 19–21). ACM, New York, pp. 91–98.

CONWAY, J. H., AND SLOANE, N. J. A. 1988. *Sphere Packings, Lattices and Groups*. Springer-Verlag, New York.

CORMEN, T. H., LEISERSON, C. E., AND RIVEST, R. L. 1990. *Introduction to Algorithms*. MIT Press, Cambridge, Mass.

DANZER, L., FONLUPT, J., AND KLEE, V. 1963. Helly's theorem and its relatives. In *Proceedings of Symposia in Pure Mathematics*, vol. 7. American Mathematical Society, 7:101–180.

DE FRAYSSEIX, H., PACH, J., AND POLLACK, R. 1988. Small sets supporting Fáry embeddings of planar graphs. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing* (Chicago, Ill., May 2–4). ACM, New York, pp. 426–433.

EDELSBRUNNER, H. 1987. *Algorithms in Combinatorial Geometry*. Springer-Verlag, New York.

EPPSTEIN, D., AND ERICKSON, J. 1994. Iterated nearest neighbors and finding minimal polytopes. *Disc. Comput. Geometry, 11,* (3), 321–350.

EPPSTEIN, D., MILLER, G. L., AND TENG, S.-H. 1993. A deterministic linear time algorithm for geometric separators and its applications. In *Proceedings of 9th Annual ACM Symposium on Computational Geometry* (San Diego, Calif., May 19–21). ACM, New York, pp. 99–108.

FÁRY, I. 1948. On straight line representing of planar graphs. *Acta. Sci. Math. 24*, 229–233.

FRIEZE, A., MILLER, G., AND TENG, S.-H. 1992. Separator-based parallel divide and conquer in computational geometry. In *Proceedings of the 4th Annual ACM Symposium on Parallel Algorithms and Architectures* (San Diego, Calif., June 29–July 1). ACM, New York, pp. 422–430.

GILBERT, J. R., MILLER, G. L., AND TENG, S.-H. 1997. Geometric mesh partitioning: Implementation and experiments. *SIAM J. Sci. Comput*., to appear.

GILBERT, J. R., HUTCHINSON, J. P., AND TARJAN, R. E. 1984. A separation theorem for graphs of bounded genus. *J. Algorithms 5*, 391–407.

GRÖTSCHEL, M., LOVÁSZ, L., AND SCHRIJVER, A. 1988. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, New York.

GUIBAS, L., PACH, J., AND SHARIR, M. 1994. Sphere-of-influence graphs in higher dimensions. In *Intuitive Geometry*, K. Boroczky and G. Fejes Toth, eds. Colloquia Mathematica Societatis J. Bolyai, vol. 63. North-Holland, Amsterdam, The Netherlands, pp. 131–137.

HAUSSLER, D., AND WELZL, E. 1987. ϵ-net and simplex range queries. *Disc. Comput. Geometry 2*, 127–151.

JORDAN, C. 1869. Sur les assemblages de lignes. *J. Reine Angew. Math. 70*, 185–190.

KABATIANSKY, G. A., AND LEVENSHTEIN, V. I. 1978. Bounds for packings on a sphere and in space. *Prob. Per. Info. 14*, 1, 3–25.

KOEBE, P. 1936. Kontaktprobleme der konformen Abbildung. *Ber. Verh. Sächs. Akademie der Wissenschaften Leipzig, Math.-Phys. Klasse 88*, 141–164.

LEIGHTON, F. T. 1983. *Complexity Issues in VLSI*. Foundations of Computing. MIT Press, Cambridge, Mass.

LEISERSON, C. E. 1983. *Area-Efficient VLSI Computation*. Foundations of Computing. MIT Press, Cambridge, Mass.

LINIAL, N., LONDON, E., AND RABINOVICH, Y. 1995. The geometry of graphs and some of its algorithmic applications. *Combinatorica 15*, 2, 215–245.

LIPTON, R. J., AND TARJAN, R. E. 1979. A separator theorem for planar graphs. *SIAM J. Appl. Math. 36*, (Apr.), 177–189.

LIPTON, R. J., ROSE, D. J., AND TARJAN, R. E. 1979. Generalized nested dissection. *SIAM J. Numer. Anal*. 16, 346–358.

MILLER, G. L., TENG, S.-H., AND VAVASIS, S. A. 1991. An unified geometric approach to graph separators. In *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science*. IEEE, New York, pp. 538–547.

MILLER, G. L., TENG, S.-H., THURSTON, W., AND VAVASIS, S. A. 1993. Automatic mesh partitioning. In *Sparse Matrix Computations: Graph Theory Issues and Algorithms*, A. George, J. Gilbert, and J. Liu, eds. *IMA Volumes in Mathematics and Its Applications*. Springer-Verlag, New York.

MILLER, G. L., TENG, S.-H., THURSTON, W., AND VAVASIS, S. A. 1997. Geometric separators for finite element meshes. *SIAM J. Sci. Comput.*, to appear.

MILLER, G. L., AND THURSTON, W. 1990. Separators in two and three dimensions. In *Proceedings of the 22th Annual ACM Symposium on Theory of Computing*, (Baltimore, Md., May 14–16). ACM, New York, pp. 300–309.

MILLER, G. L., AND VAVASIS, S. A. 1991. Density graphs and separators. In *Proceedings of the 2nd Annual ACM-SIAM Symposium on Discrete Algorithms* (San Francisco, Calif., Jan. 28–30). ACM, New York, pp. 331–336.

MOHAR, B. 1993. A polynomial time circle packing algorithm. *Disc. Comput. Geom*. 241–250.

PACH, J., AND AGARWAL, P. 1995. *Combinatorial Geometry*. Wiley-Interscience, New York.

PLOTKIN, S., RAO, S., AND SMITH, W. D. 1994. Shallow excluded minors and improved graph decomposition. In *Proceedings at the 5th Symposium Discrete Algorithms*. SIAM, New York, pp. 462–470.

PREPARATA, F. P., AND SHAMOS, M. I. 1985. *Computational Geometry: An Introduction*. Texts and Monographs in Computer Science. Springer-Verlag, New York.

SOMMERVILLE, D. M. Y. 1958. *An Introduction to the Geometry of n Dimensions*. Dover, New York.

SPIELMAN, D. A., AND TENG, S.-H. 1996a. Disk packings and planar separators. In *Proceedings of 12th Annual ACM Symposium on Computational Geometry* (Philadelphia, Pa., May 24–26). ACM, New York, pp. 349–359.

SPIELMAN, D. A., AND TENG, S.-H. 1996b. Spectral partitioning works: planar graphs and finite element meshes. UCB//CSD-96-898, U.C. Berkeley, Berkeley, California.

TENG, S.-H. 1991. Points, Spheres, and Separators: A unified geometric approach to graph partitioning. Ph.D. dissertation. CMU-CS-91-184. School of Computer Science, Carnegie-Mellon Univ., Pittsburgh, Pa.

TENG, S.-H. 1994. Combinatorial aspects of geometric graphs. MIT Tech. Rep. TR-612. May. In *Lecture Notes in Computer Science*. Springer-Verlag, New York.

THOMASSEN, C. 1980. Planarity and duality of finite and infinite graphs. *J. Combin. Theory, Series B 29*, 244–271.

TOUSSAINT, G. 1988. A graph-theoretical primal sketch. In *Computational Morphology*, G. Toussaint, ed., Elsevier-North Holland, Amsterdam, The Netherlands, pp. 229–260.

THURSTON, W. P. 1988. The geometry and topology of 3-manifolds. Princeton University Notes, Princeton, N.J.

TUTTE, W. T. 1960. Convex representations of graphs. *Proc. London Math. Soc. 10*, 3, 304–320.

TUTTE, W. T. 1963. How to draw a graph. *Proc. London Math. Soc. 13*, 3, 743–768.

UNGAR, P. 1951. A theorem on planar graphs. *J. London Math Soc. 26*, 256–262.

VALIANT, L. G. 1981. Universality consideration in VLSI circuits. *IEEE Trans. Comput. 30*, 2 (Feb.). 135–140.

VAPNIK V, N. AND CHERVONENKIS, A. YA. 1971. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Prob. Appl. 16*, 264–280.

VAVASIS, S. A. 1991. Automatic domain partitioning in three dimensions. *SIAM J. Sci. Stat. Comput. 12*, 950–970.

WYNER, A. D. 1965. Capabilities of bounded discrepancy decoding. *Bell System Tech. J. 44*, 1061–1122.

ZIEGLER, G. M. 1988. Lectures on polytopes. In *Graduate Texts in Mathematics*, vol. 152. Springer-Verlag, New York.

RECEIVED APRIL 1995; REVISED NOVEMBER 1996; ACCEPTED SEPTEMBER 1996