SIAM J. SCI. COMPUT.
Vol. 19, No. 2, pp. 364–386, March 1998

003

# GEOMETRIC SEPARATORS FOR FINITE-ELEMENT MESHES*

GARY L. MILLER†, SHANG-HUA TENG‡, WILLIAM THURSTON§, AND STEPHEN A. VAVASIS¶

**Abstract.** We propose a class of graphs that would occur naturally in finite-element and finite-difference problems and we prove a bound on separators for this class of graphs. Graphs in this class are embedded in $d$-dimensional space in a certain manner. For $d$-dimensional graphs our separator bound is $O(n^{(d-1)/d})$, which is the best possible bound. We also propose a simple randomized algorithm to find this separator in $O(n)$ time. This separator algorithm can be used to partition the mesh among processors of a parallel computer and can also be used for the nested dissection sparse elimination algorithm.

**Key words.** graph separators, finite elements, mesh partitioning, domain decomposition, computational geometry, sparse matrix computations, conformal mapping, center points

**AMS subject classifications.** 65F50, 68Q20

**PII.** S1064827594262613

## 1. Domain partitioning.

One motivation for this work is numerical solution of boundary value problems. Let $\Omega$ be an open connected region of $\mathbb{R}^d$. Suppose one is given a real-valued map $f$ on $\Omega$ and is interested in finding a map $u : \Omega \to \mathbb{R}$ such that

$$\triangle u = f \text{ on } \Omega \quad \text{and} \quad u = 0 \text{ on } \partial\Omega.$$

This problem, Poisson's equation, arises in many physical applications. Two common techniques for this problem are finite differences and finite elements. These techniques grow out of different analyses, but the end result is the same. In particular, a discrete set of nodes is inserted into $\Omega$ and a sparse system of linear equations is solved in which there is one node point and one equation for each node interior to $\Omega$. Moreover, the sparsity pattern of the system reflects interconnections of the nodes. Let the nodes and their interconnections be represented as an undirected graph $G$.

Two numerical techniques for solving this system are *domain decomposition* and *nested dissection*. Domain decomposition divides the nodes among processors of a parallel computer. An iterative method is formulated that allows each processor to

*Received by the editors January 31, 1994; accepted for publication (in revised form) April 1, 1996.

http://www.siam.org/journals/sisc/19-2/26261.html

†School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213 (glmiller@theory.cs.cmu.edu). The research of this author was supported in part by NSF grant CCR-9016641.

‡Department of Computer Science, University of Minnesota, Minneapolis, MN 55455 (steng@cs.umn.edu). The research of this author was supported in part by NSF CAREER award CCR-9502540. Part of this work was done while the author was at the Department of Mathematics and the Laboratory for Computer Science, MIT, Cambridge, MA, where the research was supported in part by Air Force Office of Scientific Research grant F49620-92-J-0125 and Advanced Research Projects Agency grant N00014-92-J-1799.

§Department of Mathematics, University of California, Berkeley, CA 94720 (wpt@math.berkeley.edu). The research of this author was supported in part by the Geometry Center in Minneapolis, the NSF Science and Technology Center for Computation and Visualization of Geometric Structures, NSF grant DMS-8920161, and by the Mathematical Sciences Research Institute, NSF grant DMS-9022140.

¶Department of Computer Science, Upson Hall, Cornell University, Ithaca, NY 14853 (vavasis@cs.cornell.edu). The research of this author was supported in part by an NSF Presidential Young Investigator award, with matching funds received from Xerox Corporation and AT&T. Part of this work was done while the author was visiting Xerox Palo Alto Research Center, Palo Alto, CA.

operate independently; see Bramble, Pasciak, and Schatz [5]. *Nested dissection*, due to George [13], George and Liu [14] and Lipton, Rose, and Tarjan [25], is a node ordering for sparse Gaussian elimination. Although originally a sequential algorithm, nested dissection also parallelizes well. For instance, Pan and Reif [37] parallelized it by writing it as sequence of matrix factors. Parallel multifrontal methods (see, e.g., [27]) are also often based on nested dissection.

For either technique it is necessary to first partition the region into subdomains. For the purpose of efficiency in both domain decomposition and nested dissection, it is important that the number of nodes in each subdomain be roughly equal, and it is also important that the size of the separator be as small as possible. For a general graph, such a decomposition may not be possible. Accordingly, it is necessary to restrict attention to classes of graphs that occur in practice in finite-difference and finite-element computations. This class is defined in the next section.

The finite-element method can be applied to many boundary value problems other than Poisson's equation; see, e.g., Johnson [22]. In most of these settings, the nested dissection and domain decomposition ideas carry over. The partitioning technique described in this paper applies to any boundary value problem posed on a spatial mesh provided the mesh satisfies quality bounds described below, and provided the pattern of nonzero entries in the discretized operator is in correspondence with the mesh topology.

First, we review the relationship between this paper and other papers by the same authors. This paper and its companion paper [31] either extend or explain several short conference papers [33, 34, 35] and one journal paper [45]. The focus of this paper is finite-element meshes; the companion paper focuses on problems arising in computational geometry.

The authors have also jointly written a survey paper [30] that surveys the results from this paper, the companion, and several additional results by various authors on efficient center point computation.

The main contributions of this paper (beyond our previous work) are

- an analysis showing that well-behaved finite-element meshes in any dimension are "overlap" graphs (defined below), and
- a complete proof of our main separator theorem showing that overlap graphs have a small separator that can be efficiently computed.

Other authors have recently looked at the mesh-partitioning problem. For instance, Pothen, Simon, and Liou [39] have a "spectral-partitioning" method based on eigenvalues of the "Laplacian" matrix of the graph. This method seems to work well in practice but does not come with any known guarantees for finite-element meshes. Hendrickson and Leland [20] have improved the spectral method with a multilevel heuristic; the improved version is much faster in practice. The Chaco package from Sandia implements this spectral method and a newer multilevel Kernighan–Lin algorithm. The original Kernighan–Lin partitioning algorithm [23] improves the partition by moving individual nodes from one subdomain to the other, and the multilevel Kernighan–Lin is able to move entire connected subgraphs [21].

Another mesh-partitioning algorithm used in practice is a graph-search heuristic due to George and Liu [14]. Leighton and Rao [24] have a partitioning method guaranteed to return a split whose separator size (see below) is within logarithmic factors of optimal, but the technique, based on flow algorithms, currently appears to be too expensive for application to large-scale meshes.

The method that we propose, unlike these previous works, assumes that the graph $G$ comes with an embedding of its nodes in $\mathbb{R}^d$. This is a very reasonable assumption

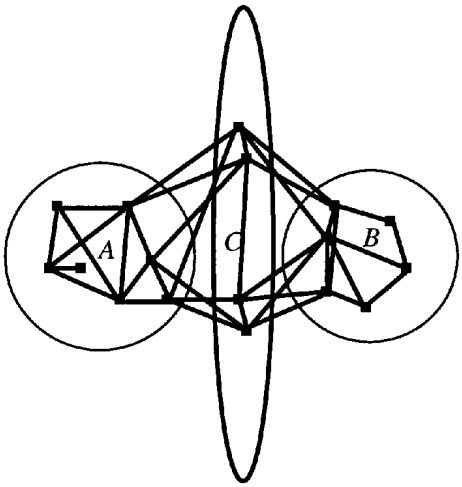


FIG. 1. *Example of $A, B, C$ in Definition* 2.1.

for finite-difference and finite-element meshes. This geometric embedding is used extensively by our algorithm. 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 bound that is the best possible bound for the class of graphs we consider.

Throughout the paper we regard the dimension $d$ as a small constant. The interesting cases for applications are commonly $d = 2$ or $d = 3$.

The remainder of the paper is organized as follows. In section 2 we introduce the concept of graph separators and define a class of graphs called overlap graphs. We also state the main theorem of the paper in section 2: there is an efficient algorithm for computing good separators of overlap graphs. In section 3 we prove that finite-element meshes belong to the class of overlap graphs. In section 4 we state our main algorithm for finding separators. The proof that this algorithm finds small separators is the focus of sections 5 through 7. In sections 8 and 9 we consider some practical issues associated with mesh partitioning.

**2. Separators and overlap graphs.** We now formally define the concept of separator.

DEFINITION 2.1. *A subset $C$ of vertices of an $n$-vertex graph $G$ is an $f(n)$-*separator *that $\delta$-splits $G$ 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 $|A|, |B| \leq \delta n$. Here, $f$ is a function and $0 < \delta < 1$.*

In this definition and for the rest of the paper, $|A|$ denotes the cardinality of a finite set $A$. The type of separator defined here is sometimes called a "vertex separator," that is, a subset $C$ of vertices of $G$ whose removal disconnects the graph into two or more graphs of smaller size; see Fig. 1. A related concept is an "edge separator," that is, a set of edges whose removal disconnects the graph. Edge separators are useful for the problem of partitioning the computational tasks of a traditional iterative algorithm (such as conjugate gradient) for a finite-element problem among the processors of a parallel computer.

Our algorithm can compute an edge separator as effectively and efficiently as it computes vertex separators. We have decided to state our theoretical bounds in terms of vertex separators, however, to be consistent with previous literature.

One of the most well known separator results is Lipton and Tarjan's result [26] that any planar graph has a $\sqrt{8n}$-separator that 2/3-splits, which improved on Ungar's [43] result. Building on this result, Gilbert, Hutchinson, and Tarjan [15] showed that all graphs with genus bounded by $g$ have an $O(\sqrt{gn})$-separator, and Alon, Seymour, and Thomas [1] proved that all graphs with an excluded minor isomorphic to the $h$-clique have an $O(h^{3/2}\sqrt{n})$-separator. These results are apparently not applicable to graphs arising as finite-element meshes when the dimension $d$ is higher than two.

The class of graphs we consider is defined by a *neighborhood* system.

DEFINITION 2.2. *Let $P = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\}$ be points in $\mathbb{R}^d$. A* $k$-ply neighborhood system *for $P$ is a set $\{B_1, \ldots, B_n\}$ of closed balls such that* (1) *$B_i$ centered at $\boldsymbol{p}_i$ and* (2) *no point $\boldsymbol{p} \in \mathbb{R}^d$ is interior to more than $k$ of $B_1, \ldots, B_n$.*

In this paper we will focus exclusively on the case that $k = 1$, i.e., the interiors of the balls are disjoint. The case when $k > 1$ is interesting for a number of geometric problems and is considered in our other paper [31].

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 $\alpha r$.

Given a neighborhood system, it is possible to define the *overlap graph* associated with the system.

DEFINITION 2.3. *Let $\alpha \geq 1$ be given, and let $\{B_1, \ldots, B_n\}$ be a* 1*-ply neighborhood system. The* $\alpha$-overlap graph *for this neighborhood system is the undirected graph with vertices $V = \{1, \ldots, n\}$ and edges*

$$E = \{(i, j) : B_i \cap (\alpha \cdot B_j) \neq \emptyset \text{ and } (\alpha \cdot B_i) \cap B_j \neq \emptyset\}.$$

The main result that we establish in this paper is as follows.

THEOREM 2.4. *Let $G$ be an $\alpha$-overlap graph, and assume $d$ is fixed. Then $G$ has an*

$$O(\alpha \cdot n^{(d-1)/d} + q(\alpha, d))$$

*separator that $(d+1)/(d+2)$ splits. A separator of the same size that $(d+1+\epsilon)/(d+2)$-splits can be computed with high probability by a randomized linear-time algorithm or randomized constant-time parallel algorithm provided that $\epsilon > 1/n^{1/2d}$.*

**3. Finite-element and finite-difference meshes.** The main result of this section is that finite-element meshes that satisfy a "shape" criterion are overlap graphs. The *finite-element method* [41] is a collection of techniques for solving boundary value problems on irregularly shaped domains. The finite-element method subdivides the domain (a subset of $\mathbb{R}^d$) into a mesh of polyhedral elements. A common choice for an element is a $d$-dimensional simplex. These simplices are arranged in a *simplicial complex*; that is, they meet only at shared subfaces. Based on the finite-element mesh, a coefficient matrix called the "assembled stiffness matrix" is defined, with variables representing unknown quantities in the mesh. Let the *finite-element graph* refer to the nonzero structure of this matrix.

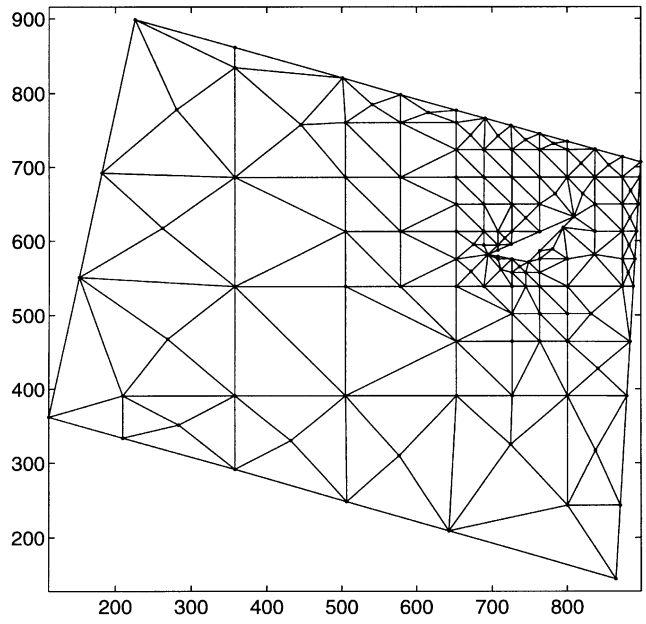


FIG. 2. *A mesh produced by a mesh generator that guarantees bounded aspect ratio.*

Associated with such a simplicial complex is its 1-*skeleton*, that is, the set of nodes accompanied by one-dimensional edges joining them. It is well known that, in the case of a piecewise-linear finite-element approximation for Poisson's equation, the nodes and edges in the finite-element graph defined in the last paragraph are in one-to-one correspondence with the nonboundary nodes and edges of the 1-skeleton of the complex. Notice that this 1-skeleton carries geometric information about the positions of the nodes and edges. In the case of higher-order elements, the finite-element graph is obtained from the 1-skeleton by introducing additional nodes interior to the faces of the simplices, and edges are introduced between every pair of nodes that share an element.

It is usually a requirement for numerical accuracy that the simplices are *well shaped* [4, 11, 46]. A common shape criterion used in mesh generation [4, 6, 36, 40] is an upper bound on the *aspect ratio* of the simplices. This term has many definitions that are all roughly equivalent [30]. One definition of aspect ratio of a simplex $T$ is the radius of the smallest sphere containing $T$ divided by the radius of the largest sphere that can be inscribed in $T$. Denote these two radii by $R(T)$ and $r(T)$ so that the aspect ratio is $R(T)/r(T)$. Figure 2 shows an example of a finite-element mesh generated by Mitchell's mesh generator and based on the algorithm in [36]. The algorithm guarantees a fixed upper bound on the aspect ratio of all triangles it produces, provided that the input polygon has no sharp angles. A three-dimensional version of that algorithm has recently been implemented by Vavasis [47].

We can show that the interior nodes in a 1-skeleton form an overlap graph. This result is easily generalized to higher-order elements. First, we prove a preliminary geometric lemma. This lemma is also used later on. For the lemma and rest of the paper, let $v_d$ be the volume of the $d$-dimensional unit ball embedded in $\mathbb{R}^d$, and let $s_{d-1}$ be the surface area of $d-1$-dimensional unit sphere embedded in $\mathbb{R}^d$. These are

well known to be

$$v_d = \frac{\pi^{d/2}}{(d/2)!}$$

for $d$ even,

$$v_d = \frac{2^{(d+1)/2}\pi^{(d-1)/2}}{1 \cdot 3 \cdot 5 \cdots d}$$

for $d$ odd, and

$$s_{d-1} = dv_d.$$

LEMMA 3.1. *Let $B$ be a ball and $S$ a sphere both embedded in $\mathbb{R}^d$. Let the radius of $B$ be $\gamma$ and the radius of $S$ be $r$ such that $r \geq \gamma$. Assume that $S$ and $0.5 \cdot B$ have a common point. Then the surface area of $S \cap B$ is at least*

$$\left(\frac{\sqrt{7}}{4}\gamma\right)^{d-1} v_{d-1}.$$

*Proof.* It suffices to prove the lemma in the special case that $\gamma = 1$ (and hence $r \geq 1$) because we can initially scale each coordinate of $\mathbb{R}^d$ by $1/\gamma$. Furthermore, without loss of generality, let $B$ be centered at the origin, and let $S$ be centered at the point $(p, 0, 0, \ldots, 0)$ with $p \geq 0$. The assumption that $S$ and $0.5 \cdot B$ have a common point means that $r - 0.5 \leq p \leq r + 0.5$. Let us consider the $(d-2)$-dimensional sphere $S'$ that is the intersection of $\partial B$ and $S$, that is, the solution to the equations

$$x_1^2 + \cdots + x_d^2 = 1,$$
$$(x_1 - p)^2 + x_2^2 + \cdots + x_d^2 = r^2.$$

Clearly, these equations have a solution if and only if

$$x_1^2 - 1 = (x_1 - p)^2 - r^2,$$

which has as its unique solution

$$x_1^* = \frac{p^2 - r^2 + 1}{2p} = \frac{p}{2} + \frac{1 - r^2}{2p}.$$

Now, let us consider the minimum and maximum possible values of $x_1^*$ over all choices of $r, p$ satisfying these constraints. First, we consider the maximum possible value. Note that both terms in the formula for $x_1^*$ are increasing as $p$ increases (because $1 - r^2 \leq 0$). Therefore, to maximize $x_1^*$ we would pick $p$ maximally to be $r + 1/2$. Substituting this in the formula for $x_1^*$ yields

$$x_{1,\max}^* = \frac{r}{2} + \frac{1}{4} + \frac{1 - r^2}{2r + 1} = \frac{1}{2} + \frac{\frac{3}{4}}{2r + 1}.$$

This is maximized for $r$ as small as possible, i.e., $r = 1$. In this case we have $x_{1,\max}^* = 3/4$.

Now, let us consider the minimum possible value of $x_1^*$. In this case, we want to pick $p = r - 1/2$, yielding

$$x_{1,\min}^* = \frac{r}{2} - \frac{1}{4} + \frac{1 - r^2}{2r - 1} = -\frac{1}{2} + \frac{\frac{3}{4}}{2r - 1}.$$

This is minimized by taking $r$ as large as possible, yielding $x^*_{1,\min} = -1/2$. Combining the upper and lower bound, we conclude that $|x^*_1| \leq 3/4$.

Now, notice that $S \cap B$, a spherical cap, is the solution to the system

$$\begin{aligned}(x_1 - p)^2 + \cdots + x_d^2 &= r^2,\\ x_1 &\leq x_1^*.\end{aligned}$$

Consider projecting this cap orthogonally onto the plane $x_1 = x_1^*$. The projection $B'$ of $S \cap B$ contains $(x_1^*, x_2, \ldots, x_d)$ if and only if

$$\begin{aligned}x_2^2 + \cdots + x_d^2 &\leq r^2 - (x_1^* - p)^2\\ &= 1 - (x_1^*)^2.\end{aligned}$$

Since orthogonal projection reduces area, the area of $S \cap B$ is at least the area of $B'$. We have already proved that $|x_1^*| \leq 3/4$, so the radius of $B'$ is at least $(1 - 9/16)^{1/2}$, i.e., at least $\sqrt{7}/4$. Thus, the area of $S \cap B$ is at least the area of $B'$, which is at least $(\sqrt{7}/4)^{d-1} v_{d-1}$. □

We now present the main theorem of this section.

THEOREM 3.2. *Let $H$ be a simplicial complex embedded in $\mathbb{R}^d$. Assume that every simplex has aspect ratio bounded by $c_1$. Let $G$ be the finite-element graph of $H$, that is, the 1-skeleton of interior nodes of $H$. Then $G$ is a subgraph of an $\alpha$-overlap graph for an $\alpha$ bounded in terms of $c_1$ and $d$.*

*Proof.* Let the interior nodes of the complex be $\{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\}$. Fix a particular $i$ and consider $\boldsymbol{p}_i$. Let $T_1, \ldots, T_q$ be the simplices adjacent to $\boldsymbol{p}_i$. Define $r_i = \min(r(T_1), \ldots, r(T_q))$. Surround node $\boldsymbol{p}_i$ with a ball $B_i$ of radius $r_i$. Carry out this construction of $r_i$ and $B_i$ for each node $\boldsymbol{p}_i$.

Note that $r_i$ is at most half the distance to the facet of $T$ opposite $\boldsymbol{p}_i$ for any simplex $T$ containing $\boldsymbol{p}_i$. Since $\boldsymbol{p}_i$ is interior to $H$, the shortest altitude to the facets of $T_1, \ldots, T_q$ opposite $\boldsymbol{p}_i$ is shorter than the distance from $\boldsymbol{p}_i$ to any other node. This shows that $B_i$ does not intersect any of the other balls $B_j$ for $j$ not equal to $i$. Thus, $B_1, \ldots, B_n$ form a 1-ply system.

Now, we prove that the edges of the 1-skeleton adjacent to $\boldsymbol{p}_i$ are covered by $\alpha \cdot B_i$. First, we argue that the number of simplices adjacent to any particular node $\boldsymbol{p}_i$ is bounded above in terms of $c_1$. The argument for this bound is as follows. Because of the aspect ratio bound, there is a lower bound on the solid angle of each simplex adjacent to $\boldsymbol{p}_i$ and therefore an upper bound $q^*$ on the number of such simplices.

Stating this argument in more detail, let $C_1, \ldots, C_q$ be the balls of radii $r(T_1), \ldots, r(T_q)$ inscribed in $T_1, \ldots, T_q$. For a particular $j$, $1 \leq j \leq q$, let $S_j$ be the sphere centered at $\boldsymbol{p}_i$ and passing through the center of $C_j$. By the lemma, the surface area of $S_j \cap C_j$ is at least $(\sqrt{7} r(T_j)/4)^{d-1} v_{d-1}$.

Now, consider the sphere $S$ of radius 1 centered at $\boldsymbol{p}_i$. For each $j$, $S_j$ is also a sphere centered at $\boldsymbol{p}_i$; hence we can expand or contract its radius to make it coincide with $S$. Let $\rho_j$ be this radius. This dilation also carries $S_j \cap C_j$ to a subregion $U_j$ of $S$. These subregions are disjoint (or perhaps have common boundary points only) for $j = 1, \ldots, q$ because $S_j \cap C_j$ and also its dilation $U_j$ lie inside the convex cone centered at $\boldsymbol{p}_i$ defined by $T_j$. The surface area of $U_j$ is at least $(\sqrt{7} r(T_j)/(4\rho_j))^{d-1} v_{d-1}$.

Note that $\rho_j \leq 2R(T_j)$ because all of $T_j$ is contained inside the sphere of radius $2R(T_j)$ centered at $\boldsymbol{p}_i$. Therefore, $S_j$ must have a smaller radius than this sphere. Thus, the surface area of $U_j$ is at least $(\sqrt{7} r(T_j)/(8R(T_j)))^{d-1} v_{d-1}$, i.e., at least $(\sqrt{7}/(8c_1))^{d-1} v_{d-1}$.

Thus, we have $q$ disjoint subsets $U_1, \ldots, U_j$ of $S$ each with surface area $(\sqrt{7}/(8c_1))^{d-1}v_{d-1}$. Since the surface area of $S$ is $s_{d-1}$, this shows that

$$q \leq \frac{s_{d-1}(8c_1)^{d-1}}{7^{(d-1)/2}v_{d-1}}.$$

Let us call this upper bound $q^*$.

Let us say that two simplices are neighbors if they share a $(d-1)$-facet. By the assumption that $\boldsymbol{p}_i$ is an interior node, we know that all of the simplices adjacent to $\boldsymbol{p}_i$ are "connected" under the transitive closure of the "neighbor" relation. Now, we claim that if two simplices $T_j, T_{j'}$ are neighbors, then $R(T_j) \geq r(T_{j'})$. This follows immediately because $r(T_{j'})$ is shorter than half the length of the shortest edge of the common face, where $R(T_j)$ is greater than half the length of the longest edge. The inequality $R(T_j) \geq r(T_{j'})$ implies that $c_1 R(T_j) \geq R(T_{j'})$.

Recall that $r_i$ is the minimum $r(T_j)$ for $j = 1, \ldots, q$; say the minimum is achieved at $j = 1$. Then the arguments in the last two paragraphs show that $R(T_j)$ for any $j$ is bounded by $c_1^{q^*} R(T_1)$. Thus, the ball of radius $2c_1^{q^*} R(T_1)$, i.e., radius $2c_1^{q^*+1} r_i$, contains $T_1, \ldots, T_q$. This ball is $(2c_1^{q^*+1}) \cdot B_i$.

This shows that $G$ is indeed a subgraph of an $\alpha$-overlap graph because all the nodes connected to $\boldsymbol{p}_i$ are vertices of $T_1, \ldots, T_q$. □

Other shape criteria weaker than an aspect ratio bound have appeared in the literature; for instance, Babuška and Aziz [2] have shown that the two-dimensional finite-element approximation converges to the true solution in the case that the largest angle of the mesh is bounded away from $\pi$ (this is a weaker condition than bounded aspect ratio). Miller, Talmor, Teng, and Walkington [29] have shown a similar result about three-dimensional Delaunay triangulations satisfying a radius-edge ratio bound. In such a triangulation, the radius of the circumscribing circle of each simplex divided by its shortest edge is bounded above by a constant, and the triangulation is a Delaunay triangulation.

In the case of a two-dimensional triangulation with an upper bound on the largest angle, the 1-skeleton of the triangulation is not necessarily an overlap graph with bounded $\alpha$. However, the bounded radius-edge Delaunay triangulation is an $\alpha$-overlap graph as argued by [29].

Another type of discretization used in solving PDEs is a finite-difference mesh with adaptive refinement; see, e.g., Fig. 3 based on a paper by Berger and Bokhari [3]. In such a mesh, it is a common rule to require that no node has neighbors more than twice as far away as its closest neighbor (otherwise the extra interpolations lead to numerical inaccuracy). It is very easy to see that such a graph is an $\alpha$-overlap graph with $\alpha = 2$.

**4. The main algorithm and stereographic projection.** We now describe our separator algorithm. Then we describe some of the details of the implementation and explain its complexity. The correctness proof of the algorithm is the subject of sections 5 through 7.

We start with two preliminary concepts. We let $\Pi$ denote the *stereographic projection* mapping from $\mathbb{R}^d$ to $S^d$, where $S^d$ is the unit $d$-sphere embedded in $\mathbb{R}^{d+1}$. Geometrically, this map may be defined as follows. Given $\boldsymbol{x} \in \mathbb{R}^d$, append "0" as the final coordinate yielding $\boldsymbol{x}' \in \mathbb{R}^{d+1}$. Then compute the intersection of $S^d$ with the line in $\mathbb{R}^{d+1}$ passing through $\boldsymbol{x}'$ and $(0, 0, \ldots, 0, 1)^T$. This intersection point is $\Pi(\boldsymbol{x})$.

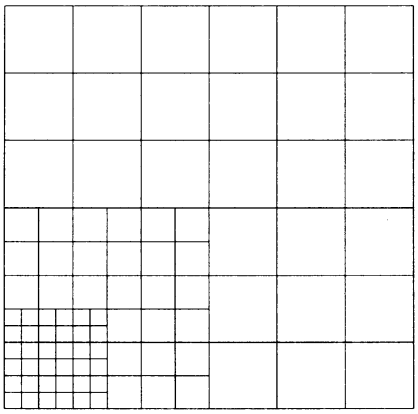

FIG. 3. *An example of finite-difference mesh with refinement.*

Algebraically, the mapping is defined as

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

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

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

where $\bar{\boldsymbol{u}}$ denotes the first $d$ entries of $\boldsymbol{u}$ and $u_{d+1}$ is the last entry. The stereographic mapping, besides being easy to compute, has a number of important properties proved below.

A second crucial concept for our algorithm is the notion of a *center point*. Given a finite subset $P \subset \mathbb{R}^d$ such that $|P| = n$, a *center point* of $P$ is defined to be a point $\boldsymbol{x} \in \mathbb{R}^d$ such that if $H$ is any open half-space whose boundary contains $\boldsymbol{x}$, then

$$|P \cap H| \leq dn/(d+1). \tag{1}$$

It can be shown from Helly's theorem [8] that a center point always exists. Note that center points are quite different from centroids. A center point is largely insensitive to "outliers" in $P$. On the other hand, a single distant outlier can cause the centroid of $P$ to be displaced by an arbitrarily large distance. In the $d = 1$ case, a center point is the same as a median for $n$ odd and is any point between the two medians for $n$ even.

*Main Separator Algorithm.*

Let $P = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\}$ be the input points in $\mathbb{R}^d$ that define the overlap graph.

1. Given $\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n$, compute $P' = \{\Pi(\boldsymbol{p}_1), \ldots, \Pi(\boldsymbol{p}_n)\}$ so that $P' \subset S^d$.
2. Compute a center point $\boldsymbol{z}$ of $P'$.
3. Compute an orthogonal $(d+1) \times (d+1)$ matrix $Q$ such that $Q\boldsymbol{z} = \boldsymbol{z}'$, where

$$\boldsymbol{z}' = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \theta \end{pmatrix}$$

such that $\theta$ is a scalar.

4. Define $P'' = QP'$ (i.e., apply $Q$ to each point in $P'$). Note that $P'' \subset S^d$, and the center point of $P''$ is $\boldsymbol{z}'$.
5. Let $D$ be the matrix $[(1-\theta)/(1+\theta)]^{1/2} I$, where $I$ is the $d \times d$ identity matrix. Let $P''' = \Pi(D\Pi^{-1}(P''))$. Below we show that the origin is a center point of $P'''$.
6. Choose a random great circle $S_0$ on $S^d$.
7. Transform $S_0$ back to a sphere $S \subset \mathbb{R}^d$ by reversing all the transformations above, i.e., $S = \Pi^{-1}(Q^{-1}\Pi(D^{-1}\Pi^{-1}(S_0)))$.
8. From $S$ compute a set of vertices of $G$ that split the graph as in Theorem 2.4. In particular, define $C$ to be vertices embedded "near" $S$, define $A$ to be vertices of $G - C$ embedded outside $S$, and define $B$ to be vertices of $G - C$ embedded inside $S$. (This step is described in section 7.)

We can immediately make the following observation: because the origin is a center point of $P'''$, and the points are split by choosing a plane through the origin, then we know that $|A| \leq (d+1)n/(d+2)$ and $|B| \leq (d+1)n/(d+2)$ regardless of the details of how $C$ is chosen. (Notice that the constant factor is $(d+1)/(d+2)$ rather than $d/(d+1)$ because the point set $P'$ lies in $\mathbb{R}^{d+1}$ rather than $\mathbb{R}^d$.) Thus, one of the claims made in Theorem 2.4 will follow as soon as we have shown that the origin is indeed a center point of $P'''$ at the end of this section.

We now provide additional details about the steps of the algorithm and also its complexity analysis. We have already defined stereographic projection used in step 1. Step 1 requires $O(nd)$ operations.

Computing a true center point in step 2 appears to a very expensive operation (involving a linear programming problem with $n^d$ constraints), but by using random (geometric) sampling, an approximate center point can be found in random constant time (independent of $n$ but exponential in $d$) [44, 19]. An *approximate* center point satisfies (1) except with $(d+1+\epsilon)n/(d+2)$ on the right-hand side, where $\epsilon > 0$ may be arbitrarily small. Alternatively, a deterministic linear-time sampling algorithm can be used in place of random sampling [28, 42], but one must again compute a center of the sample using linear programming in time exponential in $d$. See [7] and [30] for more discussion on center points and efficient algorithms for approximately computing them; see [16] for practical behavior of these randomized algorithms.

In step 3, the necessary orthogonal matrix may be represented as a single Householder reflection; see [17] for an explanation of how to pick an orthogonal matrix to zero out of all but one entry in a vector. The number of floating point operations involved is $O(d)$ independent of $n$.

In step 4 we do not actually need to compute $P''$; the set $P''$ is defined only for the purpose of analysis. Thus, step 4 does not involve computation. Note that $\boldsymbol{z}'$ is the center point of $P''$ after this transformation because when a set of points is transformed by any orthogonal transformation, a center point moves according to the same transformation (more generally, center points are similarly moved under any affine transformation). This is proved below.

In step 6 we choose a random great circle, which requires time $O(d)$. This is equivalent to choosing a plane through the origin with a randomly selected orientation. (This step of the algorithm can be made deterministic; see [10].) Step 7 is also seen to require time $O(d)$.

Finally, there are two possible alternatives for carrying out step 8, which are both described in section 7 in more detail. One alternative is that we are provided with the neighborhood system of the points (i.e., a list of $n$ balls in $\mathbb{R}^d$) as part of the input. In this case step 8 requires $O(nd)$ operations, and the test to determine which points

belong in $A$, $B$, or $C$ is a simple geometric test involving $S$. Another possibility is that we are provided with the nodes of the graph and a list of edges. In this case we determine which nodes belong in $A$, $B$, or $C$ based on scanning the adjacency list of each node, which requires time linear in the size of the graph.

We conclude this section by proving lemmas about stereographic projection and center points. These lemmas establish the claims made within the statement of the algorithm; in particular, they establish that the origin in step 5 is indeed a centerpoint of $P'''$.

LEMMA 4.1. *The mapping* $\Pi$ *is conformal.*

*Proof.* Recall that a differentiable mapping $F : \mathbb{R}^d \to \mathbb{R}^{d'}$ is said to be *conformal* if $F'(\boldsymbol{x})^T F'(\boldsymbol{x}) = \beta(\boldsymbol{x})^2 I$ for all $\boldsymbol{x}$, where $\beta(\boldsymbol{x})^2$ is a positive scalar and $I$ is the $d \times d$ identity matrix, i.e., the columns of $F'$ form an orthonormal basis multiplied by a scalar. The proof of this lemma is a straightforward computation; observe that

$$\Pi'(\boldsymbol{x}) = (2/\chi^2) \begin{bmatrix} -2\boldsymbol{x}\boldsymbol{x}^T + \chi I \\ 2\boldsymbol{x}^T \end{bmatrix},$$

and therefore

$$\begin{aligned} \Pi'(\boldsymbol{x})^T \Pi'(\boldsymbol{x}) &= (4/\chi^4)(4\boldsymbol{x}\boldsymbol{x}^T\boldsymbol{x}\boldsymbol{x}^T - 4\chi\boldsymbol{x}\boldsymbol{x}^T + \chi^2 I + 4\boldsymbol{x}\boldsymbol{x}^T) \\ &= (4/\chi^4)(4\boldsymbol{x}\boldsymbol{x}^T(\boldsymbol{x}^T\boldsymbol{x} - \chi + 1) + \chi^2 I) \\ &= (4/\chi^4)(\chi^2 I). \end{aligned}$$

To obtain the last line we used the equation $\boldsymbol{x}^T\boldsymbol{x} - \chi + 1 = 0$ by definition of $\chi$. ∎

The following lemma concerns planes in $\mathbb{R}^{d+1}$; a *plane* is defined to be the set of points $\boldsymbol{x}$ satisfying one linear equation $\boldsymbol{a}^T\boldsymbol{x} = b$, where $\boldsymbol{a}$ is a nonzero vector.

LEMMA 4.2. *Let* $V = \{\boldsymbol{y} \in \mathbb{R}^{d+1} : \boldsymbol{a}^T\boldsymbol{y} = b\}$ *be a plane in* $\mathbb{R}^{d+1}$ *that intersects* $S^d$. *The function* $\Pi^{-1}$ *maps* $S^d \cap V$ *to a sphere in* $\mathbb{R}^d$.

*Proof.* (Degenerate cases of "spheres" in $\mathbb{R}^d$ include points and planes.) Let $\boldsymbol{u}$ be a point in $S^d \cap V$. Partition $\boldsymbol{u}$ as $(\bar{\boldsymbol{u}}, u_{d+1})$, where $\bar{\boldsymbol{u}} \in \mathbb{R}^d$. Partition $\boldsymbol{a}$ in the same way. Then $\bar{\boldsymbol{a}}^T\bar{\boldsymbol{u}} = b - a_{d+1}u_{d+1}$. Assume that $\boldsymbol{u} = \Pi(\boldsymbol{x})$ for some $\boldsymbol{x} \in \mathbb{R}^d$; then we have $2\bar{\boldsymbol{a}}^T\boldsymbol{x}/\chi = b - a_{d+1}(1 - 2/\chi)$, i.e., $2\bar{\boldsymbol{a}}^T\boldsymbol{x} = (b - a_{d+1})\chi + 2a_{d+1}$, where $\chi = \boldsymbol{x}^T\boldsymbol{x} + 1$. If $b - a_{d+1} = 0$, then this set defines a plane in $\mathbb{R}^d$. Else assume $b - a_{d+1}$ is nonzero. Then the above equation may be written as

$$\boldsymbol{x}^T\boldsymbol{x} - \frac{2\bar{\boldsymbol{a}}^T\boldsymbol{x}}{b - a_{d+1}} + 1 + \frac{2a_{d+1}}{b - a_{d+1}} = 0.$$

This is the equation of a sphere in $\mathbb{R}^d$. ∎

LEMMA 4.3. *Let* $\rho$ *be the scalar in step* 5, *i.e.,* $\rho = [(1 - \theta)/(1 + \theta)]^{1/2}$, *and let* $D = \rho I$. *(Note that* $\rho$ *is well defined because* $-1 < \theta < 1$. *The center point of at least* $d + 3$ *distinct points on the sphere must be interior to the sphere itself.) Planes passing through* $\boldsymbol{z}'$ *are mapped by the transformation* $\Pi \circ D \circ \Pi^{-1}$ *to planes passing through the origin, and similarly for half-spaces.*

*Remark.* This lemma shows that the origin is the center point of $P''$. Furthermore, this lemma also shows that if an approximate center point is used in step 2 instead of an exact center point, then the transformation of step 5 also preserves approximate centering.

*Proof.* Let $V$ be a plane passing through $\boldsymbol{z}'$; such a plane has the form $V = \{\boldsymbol{u} : \bar{\boldsymbol{a}}^T\bar{\boldsymbol{u}} = a_{d+1}(\theta - u_{d+1})\}$, following the notation of the preceding lemma. (We will

prove the lemma just for the case of planes; to prove the case of half-spaces, we would instead start with $V$ defined to be $\{\boldsymbol{u} : \bar{\boldsymbol{a}}^T\bar{\boldsymbol{u}} > a_{d+1}(\theta - u_{d+1})\}$ and then carry out the same analysis.)

If we apply $\Pi^{-1}$ to such a plane, as in the last lemma, the image is

$$V' = \{\boldsymbol{x} : 2\bar{\boldsymbol{a}}^T\boldsymbol{x} = a_{d+1}(1 + \theta - \boldsymbol{x}^T\boldsymbol{x}(1-\theta))\}.$$

Now, we apply $D$; the image $DV'$ is

$$V'' = \{\boldsymbol{x} : 2\bar{\boldsymbol{a}}^T\boldsymbol{x}/\rho = a_{d+1}(1 + \theta - \boldsymbol{x}^T\boldsymbol{x}(1-\theta)/\rho^2)\}.$$

Finally, we apply $\Pi$; if the image point in $S^d$ is $(\bar{\boldsymbol{z}}, z_{d+1})$, then we know from the stereographic formulas that the preimage point in $\mathbb{R}^d$ is $\boldsymbol{x} = \bar{\boldsymbol{z}}/(1 - z_{d+1})$ and that $\boldsymbol{x}^T\boldsymbol{x} = 2/(1-z_{d+1}) - 1$. Thus, the points $\boldsymbol{z}$ in the image $V''' = \Pi \circ D \circ \Pi^{-1}(V)$ satisfy the equation

$$\frac{2\bar{\boldsymbol{a}}^T\bar{\boldsymbol{z}}}{\rho(1 - z_{d+1})} = a_{d+1}\left[1 + \theta - \frac{(2/(1 - z_{d+1}) - 1)(1-\theta)}{\rho^2}\right].$$

Multiplying through by $1 - z_{d+1}$ yields

$$\frac{2\bar{\boldsymbol{a}}^T\bar{\boldsymbol{z}}}{\rho} = a_{d+1}\left[(1+\theta)(1 - z_{d+1}) - \frac{(2 - (1 - z_{d+1}))(1-\theta)}{\rho^2}\right].$$

This equation is linear in $\boldsymbol{z}$, showing that $V'''$ is a plane in $\mathbb{R}^{d+1}$. To verify that it is a plane passing through the origin, we substitute $\bar{\boldsymbol{z}} = 0$ and $z_{d+1} = 0$ to see if we get an equation:

$$0 = a_{d+1}\left[(1+\theta) - \frac{(1-\theta)}{\rho^2}\right].$$

It is now seen that the choice $\rho = [(1-\theta)/(1+\theta)]^{1/2}$ used in step 5 does indeed make this equation hold. □

There are a few things to note about this separator algorithm.

1. For all steps except the last, the only information used about the nodes is their geometric embedding. We do not need to know their balls $B_i$'s defining the 1-ply neighborhood, we do not need to know $\alpha$, and we do not need to know explicit edges in the graph. This means that our algorithm can be applied to graphs that are suspected to be $\alpha$-overlap graphs without actually computing $\alpha$. In some circumstances (such as $k$-nearest neighbor graphs) we can construct the $k$-ply neighborhood system in linear time starting from the coordinates of the points [31].
2. For all steps except the last, a random sample of $P$ can be used in place of $P$. The size of the random sample depends on $d$; for example, for $d = 3$ we have used sample sizes of about 1200. This means that the running time of most of the algorithm is independent of $n$; see [30] and [16].

**5. Construction of a cost function.** In this section we begin the proof of the main result in Theorem 2.4, which is that $|C|$ is bounded by $O(\alpha n^{(d-1)/d} + \text{const})$ with high probability. Before starting into the details of the proof, let us provide a proof sketch.

- In this section we construct a *cost* function $g(\boldsymbol{x})$ that maps $\mathbb{R}^d$ to the nonnegative real numbers. This construction is based on the neighborhood system. The value of $g(\boldsymbol{x})$ is large if many nodes of the neighborhood system $P$ are "near" $\boldsymbol{x}$. The support of $g$ is finite.
- Most of this section is devoted to establishing the result that the integral of $g^d$ over all of $\mathbb{R}^d$ is bounded by $O(\alpha^{d/(d-1)}n)$. We call this integral the "total cost" of $g$.
- In section 6 we define the cost of a random sphere $S$ in step 6 of the algorithm by the integral of $g^{d-1}$ over the image of $S$ back in $\mathbb{R}^d$. We show that the expected value of the cost is $O(\alpha n^{(d-1)/d})$. To prove this requires the Hölder inequality, and one of the factors in the Hölder estimate is the total cost that will have been analyzed in section 5.
- Finally, in section 7 we explain which nodes should be placed in $C$ based on the neighborhood system. We then show that every node placed in $C$ (with the exception of a few nodes whose number is bounded above independently of $n$) can be "charged" against the integral defining the cost of $S$. There is a constant lower bound on the amount that each node in $C$ charges against the cost of $S$. Therefore, the total number of nodes in $C$ (besides the exceptional set) is bounded above by the cost of $S$. But we will have already shown in section 6 that the cost of $S$ is expected to be $O(\alpha n^{(d-1)/d})$, so this yields the upper bound on the expected value of $|C|$.

The intuitive reason that we can "charge" the cost of a node in $C$ against the integral of $g^{d-1}$ over $S$ is that $g^{d-1}$ gets larger wherever $S$ passes near a dense cluster of nodes. But such a cluster is precisely the place where more nodes will have to be put in $C$ to separate the graph.

As mentioned above, for the proof of Theorem 2.4 we construct a nonnegative real-valued cost function $g$ on $\mathbb{R}^d$ based on the neighborhood system. As above, let the nodes be $\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n$ and let the $\alpha$-overlap graph be defined by 1-ply neighborhood system $B_1, \ldots, B_n$. Let the radii of $B_1, \ldots, B_n$ be $r_1, \ldots, r_n$, and define $\gamma_i = 2\alpha r_i$ for $i = 1, \ldots, n$.

For each $\boldsymbol{p}_i$ we define $f_i$ as follows:

$$f_i(\boldsymbol{x}) = \begin{cases} 1/\gamma_i & \text{if } \boldsymbol{x} \in (2\alpha) \cdot B_i, \text{ i.e., } \|\boldsymbol{x} - \boldsymbol{p}_i\| \leq \gamma_i, \\ 0 & \text{otherwise.} \end{cases}$$

Notice that

$$\int_{\mathbb{R}^d} f_i^d \, dV = v_d.$$

(Recall from section 3 that $v_d$ denotes the volume of the $d$-dimensional unit sphere.) Here, and for the rest of the paper, integrations over volumes in $\mathbb{R}^d$ or $d$-dimensional surfaces in $\mathbb{R}^{d+1}$ are denoted by $dV$, and integrations over $(d-1)$-dimensional surfaces are denoted by $dS$. Next, define $f$ and $g$, nonnegative functions on $\mathbb{R}^d$, as follows:

$$f(\boldsymbol{x}) = \left( \sum_{i=1}^{n} f_i(\boldsymbol{x})^d \right)^{1/d}$$

and

$$g(\boldsymbol{x}) = \left( \sum_{i=1}^{n} f_i(\boldsymbol{x})^{d-1} \right)^{1/(d-1)}.$$

We notice immediately that

$$\int_{\mathbb{R}^d} f^d \, dV = v_d n \tag{2}$$

because this integral is equal to the sum of the integrals of the $f_i^d$.

The rest of this section is devoted to establishing an $O(n)$ upper bound on the integral of $g^d$. For this we need a series of lemmas. We use these lemmas to establish that $f^d$ and $g^d$ are always within a constant factor of each other at every point. Once this fact is established, we can then easily estimate the integral of $g^d$ since (2) is an exact formula for the integral of $f^d$.

Because the distinction between $f$ and $g$ has to do with a slight shift in the exponents, we need lemmas that establish some basic properties concerning powers of sums and sums of powers.

This first lemma is an auxiliary lemma used to prove Lemma 5.2.

LEMMA 5.1. *Let $a_1, \ldots, a_n$ be nonnegative numbers, and suppose $p \geq 1$. Then*

$$\left( \sum_{i=1}^{n} a_i \right)^p \leq p \sum_{i=1}^{n} \left[ a_i \left( \sum_{j=i}^{n} a_j \right)^{p-1} \right].$$

*Proof.* Define the function

$$\phi(x_1, \ldots, x_n) = \left( \sum_{i=1}^{n} x_i \right)^p.$$

We notice that

$$\frac{\partial \phi}{\partial x_j} = p \left( \sum_{i=1}^{n} x_i \right)^{p-1}$$

for any $j$. Let $\mathbf{a}^{(i)}$ be the vector in $\mathbb{R}^n$ given by

$$\mathbf{a}^{(i)} = (0, \ldots, 0, a_i, a_{i+1}, \ldots, a_n).$$

Then

$$\begin{aligned}
\phi(a_1, \ldots, a_n) &= \phi(\mathbf{a}^{(1)}) - \phi(\mathbf{a}^{(n+1)}) \\
&= \sum_{i=1}^{n} \left[ \phi(\mathbf{a}^{(i)}) - \phi(\mathbf{a}^{(i+1)}) \right] \\
&= \sum_{i=1}^{n} \int_0^{a_i} \frac{\partial \phi}{\partial x_i}(0, \ldots, 0, t, a_{i+1}, \ldots, a_n) \, dt \\
&= p \sum_{i=1}^{n} \int_0^{a_i} \left( t + \sum_{j=i+1}^{n} a_j \right)^{p-1} dt \\
&\leq p \sum_{i=1}^{n} a_i \cdot \left( a_i + \sum_{j=i+1}^{n} a_j \right)^{p-1} \\
&= p \sum_{i=1}^{n} a_i \cdot \left( \sum_{j=i}^{n} a_j \right)^{p-1}. \qquad \square
\end{aligned}$$

The next lemma relates two different sums of powers involving a sequence of numbers $\ldots, m_{-1}, m_0, m_1, m_2, \ldots$. Below we will express $f^d$ and $g^d$ in terms of sums of this kind.

LEMMA 5.2. *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 $\theta$ and such that at most a finite number of $m_i$'s are nonzero. Let $d \geq 2$ be an integer. Then*

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

*where $c_d$ is a positive number depending on $d$.*

*Proof.* Since at most a finite number of the $m_k$ are nonzero, then we can apply the preceding lemma because the above sums are actually finite. Applying the lemma, we see that

$$\begin{aligned}
\left(\sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)}\right)^{d/(d-1)} &\leq \frac{d}{d-1} \sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)} \cdot \left(\sum_{j=k}^{\infty} m_j 2^{-j(d-1)}\right)^{1/(d-1)} \\
&\leq \frac{d}{d-1} \sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)} \cdot \left(\sum_{j=k}^{\infty} \theta \cdot 2^{-j(d-1)}\right)^{1/(d-1)} \\
&= \frac{d}{d-1} \sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)} \cdot \left(\frac{\theta \cdot 2^{-k(d-1)}}{1 - 2^{-(d-1)}}\right)^{1/(d-1)} \\
&= c_d \theta^{1/(d-1)} \sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)} \cdot 2^{-k} \\
&= c_d \theta^{1/(d-1)} \sum_{k=-\infty}^{\infty} m_k 2^{-kd}. \qquad \square
\end{aligned}$$

The next lemma is used to establish one direction on the relation between $f^d$ and $g^d$.

LEMMA 5.3. *Let $a_1, \ldots, a_n$ be nonnegative numbers, and $d \geq 2$. Then*

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

*Proof.* See [12]. $\square$

We now come to the main result for this section, which uses the preceding lemmas.

THEOREM 5.4. *For all $\boldsymbol{x} \in \mathbb{R}^d$, the following inequalities hold:*

$$f(\boldsymbol{x})^d \leq g(\boldsymbol{x})^d \leq c'_d \alpha^{d/(d-1)} f(\boldsymbol{x})^d,$$

*where $c'_d$ is a constant depending on $d$.*

*Proof.* The first inequality follows immediately from the definitions of $f$ and $g$ and Lemma 5.3.

For the second inequality we focus on a particular point $\boldsymbol{x} \in \mathbb{R}^d$. If $f(\boldsymbol{x}) = 0$, then $g(\boldsymbol{x}) = 0$ as well, so the inequality follows. Otherwise, define for all integers $k$

$$M_k = \{i \in \{1, \ldots, n\} : 2^{-k} \leq f_i(\boldsymbol{x}) < 2^{-k+1}\}.$$

Notice that the $M_k$'s are pairwise disjoint, and their union is the set of indices $i$ such that $f_i(\boldsymbol{x}) \neq 0$.

Let $m_k$ denote the cardinality of $M_k$. We claim that $m_k \leq c''_d \alpha^d$, where $c''_d$ is a constant.

To prove this, observe that if $i \in M_k$, then $2^{-k} \leq 1/\gamma_i \leq 2^{-k+1}$, i.e., $2^{k-1} \leq \gamma_i \leq 2^k$. This means that $r_i \geq 2^{k-1}/(2\alpha)$, where $r_i$ is the radius of $B_i$. Also, since $f_i(\boldsymbol{x}) > 0$, $\|\boldsymbol{x} - \boldsymbol{p}_i\| \leq \gamma_i$, which implies $\|\boldsymbol{x} - \boldsymbol{p}_i\| \leq 2^k$. Let $B$ be the ball of radius $(1 + 1/(2\alpha))2^k$ centered at $x$. Since the ball of radius $2^k$ contains all the $\boldsymbol{p}_i$'s such that $i \in M_k$, we see that $B$ contains all the $B_i$'s for $i \in M_k$.

On the other hand, these balls have disjoint interiors because they define a 1-ply system. Accordingly, there are $m_k$ balls of radius at least $2^{k-2}/\alpha$ lying in a sphere of radius $(1 + 1/(2\alpha))2^k$, so a straightforward volume-counting argument shows

$$\begin{aligned} m_k &\leq \frac{((1 + 1/(2\alpha))2^k)^d}{(2^{k-2}/\alpha)^d} \\ &\leq \frac{(2 \cdot 2^k)^d}{(2^{k-2}/\alpha)^d} \\ &= c''_d \alpha^d. \end{aligned}$$

Now, we observe that

$$\begin{aligned} g(\boldsymbol{x})^d &= \left( \sum_{k=-\infty}^{\infty} \sum_{i \in M_k} f_i(\boldsymbol{x})^{d-1} \right)^{d/(d-1)} \\ &\leq \left( \sum_{k=-\infty}^{\infty} m_k (2^{-k+1})^{d-1} \right)^{d/(d-1)} \\ &= 2^d \left( \sum_{k=-\infty}^{\infty} m_k (2^{-k})^{d-1} \right)^{d/(d-1)} \end{aligned}$$

with $m_k$ bounded by $c''_d \alpha^d$. Now, we can apply Lemma 5.2 with the choice $\theta = c''_d \alpha^d$ to deduce that

$$g(\boldsymbol{x})^d \leq c_d 2^d (c''_d \alpha^d)^{1/(d-1)} \sum_{k=-\infty}^{\infty} m_k 2^{-kd}.$$

This summation is a lower bound on $f(\boldsymbol{x})^d$ because for each $i \in M_k$, $f_i(\boldsymbol{x})^d \geq 2^{-kd}$. This concludes the proof of the theorem. ∎

Therefore, $g^d$ is no more than a constant multiple of $f^d$, where the constant is $c'_d \alpha^{d/(d-1)}$. By (2) we have a bound of the form

$$\int_{\mathbb{R}^d} g^d \, dV \leq c_d \alpha^{d/(d-1)} n, \tag{3}$$

where $c_d$ is a different constant depending on $d$.

**6. Analysis of a random great circle.** Let $S$ be a sphere in $\mathbb{R}^d$. We define

$$\text{cost}(S) = \int_S g^{d-1} \, dS.$$

The rationale for this definition will be provided in section 7, where we prove that $\text{cost}(S)$ is proportional to the cost of separating the graph with sphere $S$, i.e., proportional to the number of graph vertices that must be removed to break all connections in $G$ from the interior of $S$ to the exterior of $S$ in step 8 of our algorithm. In this section we obtain an upper bound on the expected value of $\text{cost}(S)$ if $S$ is chosen (at random) by step 6 in our algorithm.

Recall that our separator algorithm computes a conformal mapping $F : \mathbb{R}^d \to S^d$ given by

$$F = \Pi \circ D \circ \Pi^{-1} \circ Q \circ \Pi.$$

Recall that the columns of $F'(\boldsymbol{x})$ for any $\boldsymbol{x}$ form an orthonormal basis multiplied by a nonzero scalar which we will denote $\beta(\boldsymbol{x})$. By continuity, $\beta(\boldsymbol{x})$ has the same sign for all $\boldsymbol{x} \in \mathbb{R}^d$, so without loss of generality, let us say that $\beta(\boldsymbol{x}) > 0$. Then it can be checked that for any real-valued integrable function $r$ defined on $\mathbb{R}^d$

$$\int_{\mathbb{R}^d} r(\boldsymbol{x})\, dV = \int_{S^d} r(F^{-1}(\boldsymbol{u})) \cdot \beta(F^{-1}(\boldsymbol{u}))^{-d}\, dV$$

because $\beta(\boldsymbol{x})^d$ is the determinant of $F'(\boldsymbol{x})$ when interpreted as a basis for the tangent space of $S^d$ at $F(\boldsymbol{x})$. Thus, in particular,

$$\int_{\mathbb{R}^d} g(\boldsymbol{x})^d\, dV = \int_{S^d} g(F^{-1}(\boldsymbol{u}))^d \beta(F^{-1}(\boldsymbol{u}))^{-d}\, dV, \tag{4}$$

where $g$ is the cost function from the last section. Let $h : S^d \to \mathbb{R}$ be defined as

$$h(\boldsymbol{u}) = g(F^{-1}(\boldsymbol{u}))/\beta(F^{-1}(\boldsymbol{u})).$$

Then we can conclude from (3) and (4) that

$$\int_{S^d} h(\boldsymbol{u})^d\, dV \le c'_d \alpha^{d/(d-1)} n.$$

Next, let us consider a randomly chosen great circle $S_0$ in $S^d$. The procedure for defining such a great circle is as follows. First, pick a unit-length vector $\boldsymbol{a}$ uniformly at random. Let the plane through the origin whose normal vector is $\boldsymbol{a}$ be denoted as $\boldsymbol{a}^\perp$, i.e., $\boldsymbol{a}^\perp = \{\boldsymbol{u} : \boldsymbol{a}^T \boldsymbol{u} = 0\}$. Finally, the great circle $S_0$ is $\boldsymbol{a}^\perp \cap S^d$. Thus, the set of all great circles of $S^d$ is itself parameterized by $S^d$ because $\boldsymbol{a} \in S^d$ is chosen uniformly at random.

Next, let us note that if $F(S) = S_0$ as in step 7 of our algorithm, then

$$\int_S g(\boldsymbol{u})^{d-1}\, dS = \int_{S_0} g(F^{-1}(\boldsymbol{u}))^{d-1} \beta(F^{-1}(\boldsymbol{u}))^{-(d-1)}\, dS$$

because $F$ restricted to $S$ is still a conformal mapping of one lower dimension. This shows that

$$\text{cost}(S) = \int_{\boldsymbol{u} \in S_0} h(\boldsymbol{u})^{d-1}\, dS.$$

Accordingly, we now analyze the expected value of the integral on the right-hand side of the preceding equation. This expected value is equal to

$$E[\text{cost}(S)] = \frac{1}{s_d} \int_{\boldsymbol{a} \in S^d} \int_{\boldsymbol{u} \in \boldsymbol{a}^\perp \cap S^d} h(\boldsymbol{u})^{d-1}\, dS\, dV.$$

We now interchange the order of integration; note that $\boldsymbol{u} \in \boldsymbol{a}^{\perp}$ iff $\boldsymbol{a} \in \boldsymbol{u}^{\perp}$. (To fully justify the interchange of integrals also requires an argument from differential geometry, which is in [32], concerning the volume elements in the two integrations.) We obtain

$$\begin{aligned} E[\text{cost}(S)] &= \frac{1}{s_d} \int_{\boldsymbol{u} \in S^d} \int_{\boldsymbol{a} \in \boldsymbol{u}^{\perp} \cap S^d} h(\boldsymbol{u})^{d-1} \, dS \, dV \\ &= \frac{s_{d-1}}{s_d} \int_{\boldsymbol{u} \in S^d} h(\boldsymbol{u})^{d-1} \, dV. \end{aligned} \tag{5}$$

The second line was obtained by noting that the integrand in the first line is independent of $\boldsymbol{a}$ and hence integration over $\boldsymbol{a}$ reduces to a constant factor.

Now, we apply the Hölder inequality [12]. The Hölder inequality says that for nonnegative functions $\phi$ and $\psi$ suitably integrable on a measurable set $V$ and for positive real numbers $p, q$ such that $1/p + 1/q = 1$, the following relation holds:

$$\int_V \phi\psi \leq \left( \int_V \phi^p \right)^{1/p} \cdot \left( \int_V \psi^q \right)^{1/q}.$$

We apply this inequality to our problem with $V = S^d$, $p = d/(d-1)$, $q = d$, $\phi = h^{d-1}$, and $\psi = 1$ (constant function) to obtain

$$\begin{aligned} \int_{S^d} h^{d-1} \, dV &\leq \left[ \int_{S^d} h^d \right]^{(d-1)/d} \cdot \left[ \int_{S^d} 1^d \right]^{1/d} \\ &= \left[ \int_{S^d} h^d \right]^{(d-1)/d} \cdot s_d^{1/d} \\ &\leq [c'_d \alpha^{d/(d-1)} n]^{(d-1)/d} \cdot s_d^{1/d} \\ &= c_d'^{(d-1)/d} \alpha n^{(d-1)/d} \cdot s_d^{1/d}. \end{aligned} \tag{6}$$

Combining (5) with (6) yields the following result.

THEOREM 6.1. *Let $S$ correspond to a randomly chosen great circle in the separator algorithm. Then*

$$E[\text{cost}(S)] \leq c''_d \alpha n^{(d-1)/d}.$$

Note that this is a bound on the expected value of $\text{cost}(S)$. Since $\text{cost}(S)$ is a nonnegative random variable, we know that with probability 0.5 a random trial will yield a choice of $S_1$ satisfying $\text{cost}(S_1) \leq 2E[\text{cost}(S)]$. Therefore, if we conduct, for instance, 10 random trials, and keep the best choice for $S$, then with probability exceeding 0.999 the cost will be bounded by $2E[\text{cost}(S)]$.

**7. Constructing a vertex separator from $S$.** In this section we explain how to construct a vertex separator of the overlap graph $G$ given the sphere $S$. In other words, we will partition the nodes of $G$ into $A, B, C$ to prove Theorem 2.4. For this section, we assume $S$ is a true sphere, and the degenerate case that $S$ is a plane is not analyzed. This degenerate case can be easily handled with a variant of the arguments in this section.

Recall that by the center point property (1) and Lemma 4.2 the number of vertices of $G$ strictly inside $S$ is bounded by $(d+1)n/(d+2)$, as is the number of vertices

strictly outside. (The weaker bound $(d+1+\epsilon)n/(d+2)$ holds if an approximate center point was computed in step 2 above.) For our construction, $A$ will be a subset of the vertices lying outside $S$, $B$ will be a subset of the vertices lying inside $S$, and $C$ will be vertices lying "close" to $S$. Because of this choice, we immediately establish the bounds on $|A|, |B|$ stated in Theorem 2.4.

In this section we show how to construct $C$ so that the number of vertices in $C$, other than a constant-sized exceptional set, is proportional to $\text{cost}(S)$. Since we already have established an upper bound on $\text{cost}(S)$, the argument in this section suffices to establish an upper bound on $|C|$.

Let us assume that we are given the neighborhood system $B_1, \ldots, B_n$ and the value of $\alpha$. Another possibility is that we are given edges of the graph $G$ instead; we comment on this other possibility later on.

Recall that the radii of $B_1, \ldots, B_n$ are denoted $r_1, \ldots, r_n$. Let $r$ denote the radius of $S$.

We define $C = C_1 \cup C_2$, where

$$C_1 = \{i : (\alpha \cdot B_i) \cap S \neq \emptyset \text{ and } \boldsymbol{p}_i \in \text{int}(S)\}$$

and

$$C_2 = \{i : B_i \cap S \neq \emptyset \text{ and } \boldsymbol{p}_i \in S \cup \text{ext}(S)\}.$$

LEMMA 7.1. *Let $A$ be the set of nodes of $G - C$ outside $S$, and let $B$ be the nodes of $G - C$ inside. Then $G$ has no edge between $A$ and $B$. (Note that nodes exactly on $S$ are in $C_2$ and hence in $C$.)*

*Proof.* Let $i, j$ be two nodes of $G - C$ such that $\boldsymbol{p}_i$ is inside $S$ and $\boldsymbol{p}_j$ is outside $S$. Then $(\alpha \cdot B_i) \cap S = \emptyset$ (because $i \notin C_1$); hence $\alpha \cdot B_i$ is entirely interior to $S$. Similarly, $B_j$ is entirely exterior to $S$. By definition of the overlap graph, there is no $(i, j)$ edge in $G$. ∎

We now come to the main theorem of this section, which also establishes Theorem 2.4.

THEOREM 7.2. *With this choice of $C$, $|C| \leq (4\alpha)^d + \text{cost}(S)/c_d'''$.*

*Proof.* Partition $C_1 = C_1' \cup C_1''$ and $C_2 = C_2' \cup C_2''$, where for $p = 1, 2$ we define

$$C_p' = \{i \in C_p : 2\alpha r_i > r\}$$

and

$$C_p'' = \{i \in C_p : 2\alpha r_i \leq r\}.$$

We bound $|C_1' \cup C_2'|$ and $|C_1'' \cup C_2''|$ separately.

First, we analyze $C_1' \cup C_2'$, which we write as $C'$. Note that for every $i \in C'$, $r_i > r/(2\alpha)$. We replace each ball $B_i$ for $i \in C'$ with a smaller ball $B_i'$ of radius exactly $r/(2\alpha)$ such that $B_i' \subset B_i$. For $i \in C_1'$ we simply define $B_i' = r/(2\alpha r_i) \cdot B_i$ to obtain this result. For $i \in C_2'$ we shrink the radius of $B_i$ by factor $r/(2\alpha r_i)$, and also we displace the center so as to maintain the property that $B_i' \cap S \neq \emptyset$.

Let $n_1 = |C'|$. Observe that we have constructed $n_1$ balls of radius exactly $r/(2\alpha)$, all of whose centers are within distance $r + r/(2\alpha)$ of the center of $S$. This means that all of these balls are contained in a ball of radius $2r$ centered at the center of $S$. Note also that these balls $B_i'$ for $i \in C'$ are pairwise disjoint because the original $B_i$'s have disjoint interiors. Therefore, a volume argument shows that

$$n_1 \leq \frac{(2r)^d}{(r/(2\alpha))^d} = (4\alpha)^d.$$

Next, we examine $C_1'' \cup C_2''$, which we write as $C''$. Fix a particular $i \in C''$. Notice that $B_i$ has radius at most $r/(2\alpha)$. Let $B_i' = (2\alpha) \cdot B_i$, so that $B_i'$ has radius $\gamma_i$ (recall that $\gamma_i$ was defined in section 5). Observe that, by definition of $C_1$ and $C_2$ above, we are guaranteed that $(0.5 \cdot B_i') \cap S$ (which is the same as $(\alpha \cdot B_i) \cap S$) is nonempty. Furthermore, by construction of $C''$ we know that $\gamma_i \leq r$. Therefore, we can apply Lemma 3.1 to conclude that the area of $B_i' \cap S$ is at least $(\sqrt{7}\gamma_i/4)^{d-1} v_{d-1}$. Note that the value of $f_i$ on $B_i'$ is precisely $1/\gamma_i$. Therefore,

$$\begin{aligned}
\int_S f_i^{d-1} &= \int_{B_i' \cap S} f_i^{d-1} \\
&= \frac{1}{\gamma_i^{d-1}} \cdot \text{area}(B_i' \cap S) \\
&\geq \frac{1}{\gamma_i^{d-1}} \cdot \left( \frac{\sqrt{7}\gamma_i}{4} \right)^{d-1} v_{d-1} \\
&= c_d'''.
\end{aligned}$$

Here, $c_d'''$ is a positive constant depending only on $d$.

Therefore,

$$\begin{aligned}
\text{cost}(S) &= \int_S g^{d-1}\, dS \\
&= \int_S \sum_{i=1}^{n} f_i^{d-1}\, dS \\
&= \sum_{i=1}^{n} \int_S f_i^{d-1}\, dS \\
&\geq \sum_{i \in C''} \int_S f_i^{d-1}\, dS \\
&\geq |C''| \cdot c_d'''.
\end{aligned}$$

Combining the upper bounds on $|C'|$ and $|C''|$ proves the theorem. □

We have shown how to carry out step 8 of the main algorithm, namely, deducing a vertex separator $C$ from the sphere $S$. We have also established the bound on $C$. The construction so far seems to require explicit knowledge of $B_1, \ldots, B_n$ and of $\alpha$. If we are not given these items as part of the input, but instead we have edges of $G$ represented explicitly, then clearly we can disconnect $G$ into two pieces by removing all edges connecting the interior of $S$ to the exterior. In this paper we are focusing on vertex separators, so we must find a set of vertices that disconnects $G$. If we let $E_1$ be the set of edges passing through $S$, then we can find a vertex separator $C$ by arbitrarily taking one endpoint of every edge in $E_1$. For the special case of overlap graphs arising from finite-element methods, which have bounded degree, it can be shown that this simple heuristic does indeed produce a vertex set $C$ with the bound stated in Theorem 2.4. Alternatively, we can find in polynomial time the minimum set of vertices $C$ of $G$ that "cover" $E_1$, where "cover" means that at least one endpoint of every edge in $E_1$ is in $C$; see [9] or [38] for this algorithm.

It can be shown that $\alpha n^{(d-1)/d}$ is the best possible bound on a separator set for an $n$-vertex $\alpha$-overlap graph. Define graph $G$, whose nodes are an $m \times m \times \cdots \times m$ array of nodes arranged in a $d$-dimensional unit-spaced lattice (so that $n = m^d$) and whose

edges connect all neighbors within distance $\alpha$. This is clearly seen to be an overlap graph, and Vavasis [45] shows that any partitioning of this graph into constant-sized pieces must involve a separator set of at least const $\cdot\, \alpha n^{(d-1)/d}$ nodes.

**8. Practical issues.** Let us call the algorithm defined in section 4 Weak-Split; it produces a partition in which the ratio of the size of the larger of $G_1$, $G_2$ to the smaller is at most $d+1+o(1)$. In practice, one often wants a split in which $G_1$ and $G_2$ have no more than half the nodes, i.e., a ratio of $1+o(1)$. (Experimental results [16] suggest that for practical examples, our algorithm produces approximately a 45%–55% split after 10 trials for the $d=3$ case, much better than the worst case 20%–80% split claimed in Theorem 2.4.) There is a standard technique originally due to [26] that derives an algorithm Strong-Split using Weak-Split as a subroutine. Strong-Split yields an even split at the expense of a greater running time (by a constant factor) and a larger constant factor in the bound on the size of the separator.

Another practical issue is a splitting into more than one subdomain. If the number of domains desired $p$ is a power of 2, this is accomplished by applying Strong-Split recursively to get domains of the desired size. The total separator size in this case is $O(p^{1/d}\alpha n^{(d-1)/d})$. This approach can be generalized for a number of subdomains $p$ not a power of 2.

**9. Conclusions and open questions.** An important question is, given a graph without an embedding, can its nodes be embedded in $\mathbb{R}^d$ to make it a subgraph of an overlap graph?

Also, fast deterministic algorithms for computing approximate center points would be very useful. Deterministic linear-time approximate algorithms are known but are not efficient enough for practical use.

It is also interesting to determine how our algorithm performs in practice compared with other current algorithms, such as the spectral method. This is the subject of recent work by Gilbert, Miller, and Teng [16], who also propose some additional heuristics not described here. The results of [16] can be summarized as follows. Typically, 30 random choices for the sphere separator sufficed. The spectral method used by [16] came from [20], where heuristic local improvement is used as well as spectral partitioning. Our geometric separation algorithm usually ran faster than spectral partitioning (but in both cases there are many possible heuristics that could speed up either one). The quality of the separators (in terms of the balance and the size of the cut) from the geometric method were about the same as from the spectral method for most test cases. In some cases spectral did better; in others, geometric did better. Though the spectral method compares favorably with other heuristics, Guattery and Miller [18] recently found a class of graphs on which the spectral method performs poorly.

**Acknowledgments.** We would like to thank David Applegate, Marshall Bern, David Eppstein, John Gilbert, Bruce Hendrickson, Ravi Kannan, Michael Klugerman, Tom Leighton, Mike Luby, Oded Schramm, Doug Tygar, and Kim Wagner for invaluable help and discussions.

REFERENCES

[1] N. Alon, P. Seymour, and R. Thomas, *A separator theorem for graphs with an excluded minor and its applications*, in Proc. of the 22th Annual ACM Symposium on Theory of Computing, ACM, New York, 1990, pp. 293–299.

[2] I. Babuška and A. K. Aziz, *On the angle condition in the finite element method*, SIAM J. Numer. Anal., 13 (1976), pp. 214–226.
[3] M. J. Berger and S. Bokhari, *A partitioning strategy for nonuniform problems on multiprocessors*, IEEE Trans. Comput., C-36 (1987), pp. 570–580.
[4] M. Bern, D. Eppstein, and J. R. Gilbert, *Provably good mesh generation*, in Proc. 31st Annual Symposium on Foundations of Computer Science, IEEE, Piscataway, NJ, 1990, pp. 231–241.
[5] J. H. Bramble, J. E. Pasciak, and A. H. Schatz, *An iterative method for elliptic problems on regions partitioned into substructures*, Math. Comp., 46 (1986), pp. 361–9.
[6] L. P. Chew, *Guaranteed Quality Triangular Meshes*, Tech. report 89–893, Department of Computer Science, Cornell University, Ithaca, NY, 1989.
[7] K. Clarkson, D. Eppstein, G. L. Miller, C. Sturtivant, and S.-H. Teng, *Approximating center points with iterated radon points*, Internat. J. Comput. Geom. Appl., 6 (1996), pp. 357–377.
[8] L. Danzer, J. Fonlupt, and V. Klee, *Helly's theorem and its relatives*, in Proc. of Symposia in Pure Mathematics, American Mathematical Society, Providence, RI, 7 (1963), pp. 101–180.
[9] A. L. Dulmage and N. S. Mendelsohn, *Coverings of bipartite graphs*, Canadian J. Math., 10 (1958), pp. 517–534.
[10] D. Eppstein, G. L. Miller, and S.-H. Teng, *A deterministic linear time algorithm for geometric separators and its applications*, Fund. Inform., 22 (1995), pp. 309–330.
[11] I. Fried, *Condition of finite element matrices generated from nonuniform meshes*, AIAA J., 10 (1972), pp. 219–221.
[12] J. E. L. G. Hardy and G. Pólya, *Inequalities*, 2nd ed., Cambridge University Press, Cambridge, 1952.
[13] J. A. George, *Nested dissection of a regular finite element mesh*, SIAM J. Numer. Anal., 10 (1973), pp. 345–363.
[14] J. A. George and J. W. H. Liu, *An automatic nested dissection algorithm for irregular finite element problems*, SIAM J. Numer. Anal., 15 (1978), pp. 1053–1069.
[15] J. Gilbert, J. Hutchinson, and R. Tarjan, *A separation theorem for graphs of bounded genus*, J. Algorithms, 5 (1984), pp. 391–407.
[16] J. Gilbert, G. Miller, and S.-H. Teng, *Geometric mesh partitioning: implementation and experiments*, in Proc. of the 9th International Parallel Processing Symposium, IEEE, Piscataway, NJ, 1995, pp. 418–427; also Tech. report CSL-94-13, Xerox Palo Alto Research Center, Palo Alto, CA; SIAM J. Sci. Comput., to appear.
[17] G. H. Golub and C. F. V. Loan, *Matrix Computations*, 2nd ed., The Johns Hopkins University Press, Baltimore, MD, 1989.
[18] S. Guattery and G. L. Miller, *On spectral partitioning methods*, in Proc. of the 6th ACM–SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, PA, 1995, pp. 233–242.
[19] D. Haussler and E. Welzl, *$\epsilon$-net and simplex range queries*, Discrete Comput. Geom., 2 (1987), pp. 127–151.
[20] B. Hendrickson and R. Leland, *An Improved Spectral Graph Partitioning Algorithm for Mapping Parallel Computations*, SIAM J. Sci. Comput., 16 (1995), pp. 452–469.
[21] B. Hendrickson and R. Leland, *A multilevel algorithm for partitioning graphs*, in Proc. Supercomputing '95, ACM, New York, 1995.
[22] C. Johnson, *Numerical Solution of Partial Differential Equations by the Finite Element Method*, Cambridge University Press, Cambridge, 1987.
[23] B. Kernighan and S. Lin, *An efficient heuristic procedure for partitioning graphs*, Bell System Technical Journal, 29 (1970), pp. 291–307.
[24] F. T. Leighton and S. Rao, *An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms*, in Proc. 29th Annual Symposium on Foundations of Computer Science, IEEE, Piscataway, NJ, 1988, pp. 422–431.
[25] R. J. Lipton, D. J. Rose, and R. E. Tarjan, *Generalized nested dissection*, SIAM J. Numer. Anal., 16 (1979), pp. 346–358.
[26] R. J. Lipton and R. E. Tarjan, *A separator theorem for planar graphs*, SIAM J. Appl. Math., 36 (1979), pp. 177–189.
[27] J. W. H. Liu, *The multifrontal method for sparse matrix solution: theory and practice*, SIAM Rev., 34 (1992), pp. 82–109.
[28] J. Matoušek, *Approximations and optimal geometric divide-and-conquer*, in Proc. 23rd ACM Symposium Theory of Computing, ACM, New York, 1991, pp. 512–522.

[29] G. L. Miller, D. Talmor, S.-H. Teng, and N. Walkington, *A Delaunay based numerical method for three dimensions: generation, formulation, and partition*, in Proc. of the 27th Annual ACM Symposium on the Theory of Computing, ACM, New York, 1995, pp. 683–692.

[30] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis, *Automatic mesh partitioning*, in Graph Theory and Sparse Matrix Computation, A. George, J. Gilbert, and J. Liu, eds., Vol. 56 of IMA Vol. Math. Appl., Springer, New York, 1993, pp. 57–84.

[31] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis, *Separators for sphere-packings and nearest neighborhood graphs*, J. ACM, 44 (1997), pp. 1–29.

[32] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis, *Geometric Separators for Finite Element Meshes: Appendix Concerning Volume Elements*, 1994, unpublished; available via ftp://ftp.cs.cornell.edu/pub/vavasis/papers/diffgeo.ps.

[33] G. L. Miller, S.-H. Teng, and S. A. Vavasis, *A unified geometric approach to graph separators*, in Proc. 31nd Annual Symposium on Foundations of Computer Science, IEEE, Piscataway, NJ, 1991, pp. 538–547.

[34] G. L. Miller and W. Thurston, *Separators in two and three dimensions*, in Proc. of the 22th Annual ACM Symposium on Theory of Computing, Baltimore, May 1990, ACM, New York, pp. 300–309.

[35] G. L. Miller and S. A. Vavasis, *Density graphs and separators*, in Proc. of the ACM–SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, PA, 1991, pp. 331–336.

[36] S. A. Mitchell and S. A. Vavasis, *Quality mesh generation in three dimensions*, in Proc. of the ACM Computational Geometry Conference, 1992, ACM, New York, pp. 212–221; also Tech. report C.S. TR 92-1267, Cornell University, Ithaca, NY.

[37] V. Pan and J. Reif, *Efficient parallel solution of linear systems*, in Proc. of the 17th Annual ACM Symposium on Theory of Computing, ACM, New York, 1985, pp. 143–152.

[38] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*, Prentice–Hall, Englewood Cliffs, NJ, 1982.

[39] A. Pothen, H. D. Simon, and K.-P. Liou, *Partitioning sparse matrices with eigenvectors of graphs*, SIAM J. Matrix Anal. Appl., 11 (1990), pp. 430–452.

[40] J. Ruppert, *A new and simple algorithm for quality* 2*-dimensional mesh generation*, in Proc. 4th Symp. Discrete Algorithms, SIAM, Philadelphia, PA, 1993, pp. 83–92.

[41] G. Strang and G. J. Fix, *An Analysis of the Finite Element Method*, Prentice–Hall, Englewood Cliffs, NJ, 1973.

[42] S.-H. Teng, *Points, Spheres, and Separators: A Unified Geometric Approach to Graph Partitioning*, Ph.D. thesis, CMU-CS-91-184, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1991.

[43] P. Ungar, *A theorem on planar graphs*, J. London Math. Soc., 26 (1951), pp. 256–262.

[44] V. N. Vapnik and A. Y. Chervonenkis, *On the uniform convergence of relative frequencies of events to their probabilities*, Theory Probab. Appl., 16 (1971), pp. 264–280.

[45] S. A. Vavasis, *Automatic domain partitioning in three dimensions*, SIAM J. Sci. Statist. Comput., 12 (1991), pp. 950–970.

[46] S. A. Vavasis, *Stable finite elements for problems with wild coefficients*, SIAM J. Numer. Anal., 33 (1996), pp. 890–916.

[47] S. A. Vavasis, *QMG Version* 1.0, available via ftp://ftp.cs.cornell.edu/pub/vavasis/qmg1.0.tar.gz and http://www.cs.cornell.edu/home/vavasis/qmg-home.html (1995).