Density Graphs and Separators

Gary L. Miller*
Stephen A. Vavasis**

TR 90-1169
November 1990

Department of Computer Science
Cornell University
Ithaca, NY 14853-7501

*Supported in part by the National Science Foundation under grant DCR-8713489.
**Supported in part by the U.S. Army Research Office through the Mathematical Sciences Institute of Cornell University.

Density graphs and separators

Gary L. Miller*
School of Computer Science
Carnegie Mellon University &
Department of Computer Science
University of Southern California

Stephen A. Vavasis†
Department of Computer Science
Cornell University

November 7, 1990

Abstract

We propose a class of graphs that would occur naturally in finite-element problems, and we prove a bound on separators for this class of graphs. For three-dimensional graphs, our separator bound is $O(N^{2/3})$. We also propose a simple randomized algorithm to find this separator in $O(N)$ time. Such an algorithm would be used as a preprocessing step for the domain decomposition method of efficiently solving a finite-element problem on a parallel computer.

This paper generalizes "local graphs" of Vavasis [1990] to the case of graphs with varying densities of nodes. It also generalizes aspects of Miller and Thurston's [1990] "stable graphs."

1 Separators and domain partitioning

Motivation for this work is Poisson's equation. Let Ω be an open connected region of $\mathbb{R}^3$. Suppose one is given a real-valued map f on Ω, and is

*Supported in part by National Science Foundation grant DCR-8713489.

†Supported in part by the U.S. Army Research Office through the Mathematical Sciences Institute of Cornell University.

interested in finding a map $u : \Omega \to \mathbb{R}$ such that

$$\begin{aligned} \triangle u &= f \text{ on } \Omega, \text{ and} \\ u &= 0 \text{ on } \partial\Omega. \end{aligned}$$

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 Ω and a sparse system of linear equations is solved in which there is one node point and one equation for each node interior to Ω. 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 operate independently. See Bramble, Pasciak and Schatz [1986], and Chan and Resasco [1987], and Bjørstad and Widlund [1986]. *Nested dissection,* due to George [1973], George and Liu [1978] and Lipton, Rose and Tarjan [1979], is a node ordering for sparse Gaussian elimination. Although originally a sequential algorithm, nested dissection also parallelizes well as shown by Pan and Reif [1985a,b].

For either technique it is necessary to first partition the region into subdomains. This is the goal of the paper at hand. The partitioning is accomplished by partitioning the nodes into $p+1$ disjoint subsets, say $G_1, \ldots, G_p, G_\sim$. Sets $G_1, \ldots, G_p$ are the *subdomains*, and set $G_\sim$ is the *boundary* or *separator.* It is required that no edge connect G_i to G_j if $1 \leq i < j \leq p$; all paths between distinct subdomains must go through the boundary.

For the purpose of efficiency in a domain decomposition algorithm, it is important for the number of nodes in each set G_i to be roughly equal, and it is also important for the size of $G_\sim$ to be as small as possible. In general, such a decomposition may not be possible; see the counterexamples in Vavasis [1990] or Miller and Thurston [1990]. Accordingly, it is necessary to restrict attention to classes of graphs that occur in practice in finite element computations. We propose the following definition.

Definition 1 *Let G be an undirected graph and let π be an embedding of its nodes in $\mathbb{R}^d$. Then we say that π is an embedding of* **density** *α if the*

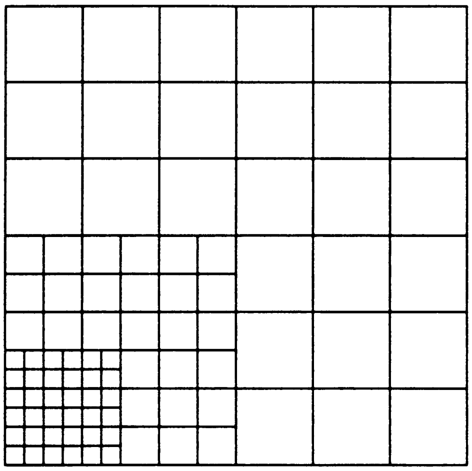

Figure 1: An example of a graph considered by Berger and Bokhari.

following inequality holds for all vertices v in G. Let u be the closest node to v. Let w be the farthest node from v that is connected to v by an edge. Then

$$\frac{\|\pi(w)-\pi(v)\|}{\|\pi(u)-\pi(v)\|} \leq \alpha.$$

In general, G is a **density graph** *if there exist a π and $\alpha > 0$ such that π is an embedding of density α*

Here and elsewhere in the paper, the norms are Euclidean norms. The idea of this definition is that a node can be connected only to nodes in its immediate neighborhood. This type of graph arises often in finite differences and finite elements; see for example Figure 1 based on Berger and Bokhari [1987] or Figure 2 generated by Chew's [1989] mesh generator. The figures depict graphs embedded in $\mathbb{R}^2$. We make further remarks about this definition in the next section. The importance of this definition is as follows. We are able to show that any N-node density graph G has a partition into $G_1, G_2, G_\sim$ such that $G_\sim$ has at most $cN^{(d-1)/d}$ nodes and such that G_1 and G_2 are no more than a constant fraction of the original domain. Here, c is a constant that depends on α and d in Definition 1. This result will take

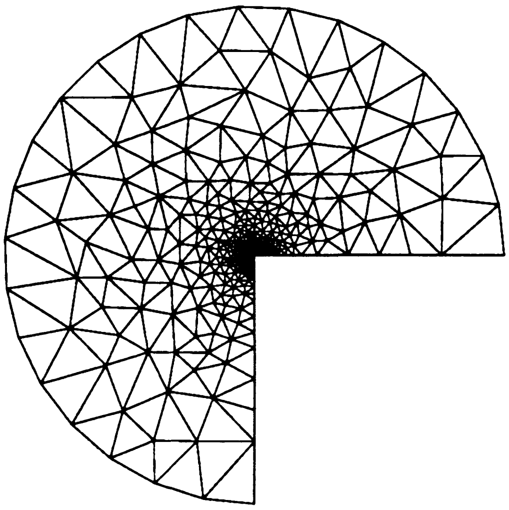

Figure 2: A graph generated by Chew's mesh generator

care of the $p = 2$ case, and partitions into larger numbers of subdomains can be accomplished recursively.

This paper allows d to be arbitrary, although the most interesting cases for numerical analysis are $d = 2$ and $d = 3$, in which case the separator is bounded by $O(N^{1/2})$ and $O(N^{2/3})$ respectively.

We remark that graph partitioning has been extensively studied for graphs more general than meshes. The literature is too extensive to survey here, but we mention work by Pothen, Simon and Liu [1989] which contains an algorithm for this problem as well as a bibliography of earlier work.

2 Comparison to other classes of graphs

The point of Definition 1 is that no node has an edge connecting it to a node very distant from its own neighborhood. We make the following trivial observation: if G has any edges, then the parameter α in the definition must be at least 1. In addition, one can easily show (see Vavasis [1990]) that any N-node graph is a density graph if we allow α as large as $cN^{1/d}$. Accordingly, the interesting case is when α is bounded independently of N.

Such behavior is expected from triangulations generated by automatic mesh generators. For example, Chew [1989] has a two-dimensional mesh generator in which all triangles have angles no less than 30° and no more than 120°. A mesh generator by Bern, Eppstein and Gilbert [1990] has a similar property.

One can easily prove that such a graph satisfies the $d = 2$ case of Definition 1 provided that assumptions are made about the shape of the boundary of the domain (see below). Similarly, a finite-differences mesh with mesh refinement will also satisfy such a condition provided that no more than one level of refinement is done per cell (as in Figure 1).

Definition 1 is a strict generalization of "local graphs" defined by Vavasis [1990]. In particular, that paper assumed that there was an upper bound on the ratio of longest edge in the whole graph to the smallest node separation in the graph. This means that such a definition could not handle graphs like the two figures in which the density of the elements varies from one region of the domain to the other. Like our main theorem, Vavasis had an $O(N^{(d-1)/d})$ bound on the separator.

Definition 1 is a partial generalization of the Miller and Thurston's class of "stable" graphs. A stable graph must have edges corresponding to the edges of a triangulation, and there must be a lower bound on the aspect ratio of each tetrahedron in the triangulation. "Aspect ratio" refers ratio of the inscribed sphere diameter to the circumscribing sphere diameter of any tetrahedron in the triangulation. Our class of density graphs do not have to be triangulations.

Moreover, there are certain kinds of triangulations that fit the density definition but violate the aspect ratio definition. For example, the tetrahedron formed by four coplanar points arranged in a square has aspect ratio of zero but would not violate the density-graph condition.

Density graphs are not a generalization of stable graphs. There are examples of stable graphs that are not density graphs because the aspect-ratio condition does not require "external boundary" nodes to be well-separated. The concept of external boundary nodes is well defined in the case of a triangulation but does not have a meaning for density graphs of this paper.

Figure 3 shows an example of a stable graph in two dimensions, that, for the embedding depicted, would not be a density graph for $\alpha < 20$ because vertices x and y are very close together. The fact that external boundary

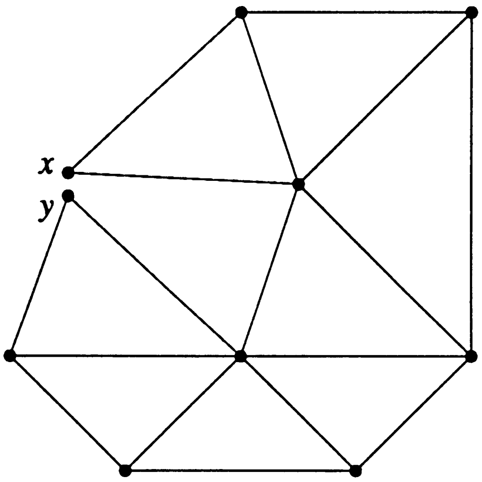

Figure 3: A stable graph with two nearby nodes.

nodes of stable graphs can be close apparently requires an extra term in the estimated size of the separator. Miller and Thurston are able to prove a bound on the separator size of the form $T_1 + T_2$, where T_1 is $O(N^{(d-1)/d})$ (the same bound in the present paper) and T_2 depends linearly on the the number of boundary nodes. Accordingly, the bound on the separator size in this paper is better, but this is because our definition excludes a troublesome case for Miller and Thurston's results.

3 A function g based on the graph G.

The main theorem of this paper is as follows.

Theorem 1 *Let G be an N-node density graph embedded in $\mathbb{R}^d$ with parameter α. Then in randomized $O(N)$ time and also in randomized NC one can find a partition $(G_1, G_2, G_\sim)$ of G such that G_1 and G_2 contain at most $(d+1)N/(d+2)$ nodes and such that $G_\sim$ contains at most $c\alpha N^{(d-1)/d} + c'$ nodes, where c depends on d, and c' depends on α and d.*

For the proof of the theorem we will construct a real-valued function g on $\mathbb{R}^d$ based on the graph. We will find a separator for the function, and deduce the existence of a graph separator. This technique of constructing a function to model the node density is taken from Miller and Thurston, although our specific construction is different. The proof of the theorem is spread over three sections.

We start with a graph G with N nodes as in the theorem. We define functions f and g as follows. Assume the nodes of G are numbered $1, \ldots, N$. Let $\pi(i)$ be the point in $\mathbb{R}^d$ where node i is embedded. Let ρ_i be the distance from node i to its most distant graph neighbor. For the ith node in G we define a function f_i as follows:

$$f_i(x) = \begin{cases} 1/\rho_i & \text{if } \|x - \pi(i)\| \le \rho_i, \text{ or} \\ 0 & \text{otherwise.} \end{cases}$$

Notice that

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

where v_d is the volume of the d-dimensional unit sphere. Here and in the rest of the paper, integrations over volumes in $\mathbb{R}^d$ are denoted with dV, and integrations over $d-1$-dimensional surfaces in $\mathbb{R}^d$ are denoted by dA. Next, define f and g pointwise as follows:

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

and

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

We notice immediately that

$$\int_{\mathbf{R}^d} f^d \, dV = v_d N \tag{1}$$

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

We would like to establish an $O(N)$ upper bound on the integral of g^d. This is the purpose of the upcoming series of lemmas. We remark that this proof would be easier if we were willing to accept poorer bounds in terms of α. This first group of lemmas establishes some basic properties concerning powers of sums and sums of powers.

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

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

Proof. Define the function

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

We notice that

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

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

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

∎

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

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

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

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

∎

Lemma 3 *Let $\phi(t)$ be a convex real-valued function defined on the nonnegative numbers such that $\phi(0) = 0$, and let $a_1, \ldots, a_n$ be a sequence of nonnegative numbers. Then*

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

Proof. For the $n = 2$ case, we observe by definition of convexity that

$$\frac{a_1}{a_1 + a_2}\phi(a_1 + a_2) + \frac{a_2}{a_1 + a_2}\phi(0) \geq \phi(a_1).$$

The second term on the left drops out by assumption. Similarly,

$$\frac{a_2}{a_1 + a_2}\phi(a_1 + a_2) \geq \phi(a_2).$$

Adding these inequalities gives the $n = 2$ result. The general case is now proved by induction. ∎

Lemma 4 *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. Let $\phi(t)$ be the function $t \mapsto t^{d/(d-1)}$. Then ϕ is a convex function on the nonnegative numbers. We do the following calculation using the previous lemma.

$$\begin{aligned}
\left(\sum_{i=1}^{n} a_i^d\right)^{1/d} &= \left(\sum_{i=1}^{n} \phi(a_i^{d-1})\right)^{1/d} \\
&\leq \left[\phi\left(\sum_{i=1}^{n} a_i^{d-1}\right)\right]^{1/d} \\
&\leq \left[\sum_{i=1}^{n} a_i^{d-1}\right]^{1/(d-1)}.
\end{aligned}$$

∎

We now come to the main result for this section. This result immediately leads to a bound on the integral of g^d over $\mathbb{R}^d$. In this result we use the preceding lemmas.

Lemma 5 *For all $x \in \mathbb{R}^d$, the following inequalities hold:*

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

where c'_d is a constant depending on d.

Proof. The first inequality follows immediately from the definitions of f and g and Lemma 4.

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

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

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

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, we rely on the density graph condition. Observe that $\pi(i)$ for $i \in M_k$ must lie within distance 2^k of x by definition of f_i. On the other hand, the longest edge adjacent to $\pi(i)$ is of length at least 2^{k-1}, also by definition of f_i. Surround each point i for $i \in M_k$ with a ball of radius $2^{k-2}/\alpha$. These balls must be disjoint for the following reason. The points $\pi(i)$ for $i \in M_k$ must be distance at least $2^{k-1}/\alpha$ from one another by definition of a density graph. Accordingly, there are m_k disjoint balls of radius at least $2^{k-2}/\alpha$ around these m_k points. All of these balls lie in a sphere of radius $2^k + 2^{k-2}/\alpha$, that is, a sphere of radius $1.25 \cdot 2^k$ (since $\alpha \geq 1$) since the centers of the balls are within 2^k.

The volume of each of the m_k smaller balls is $2^{(k-2)d} v_d / \alpha^d$, and the volume of the enclosing sphere is at most $1.25^d v_d \cdot 2^{kd}$. This gives an upper bound on m_k because the balls are disjoint. In particular, m_k is bounded by the quotient of these quantities, i.e.

$$\begin{aligned} m_k &\leq \frac{1.25^d 2^{kd}}{2^{(k-2)d}/\alpha^d} \\ &\leq c_d'' \alpha^d \end{aligned}$$

Now we observe that

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

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

with m_k bounded by $c''_d \alpha^d$. Now we can apply Lemma 2 with the choice $\theta = c''_d \alpha^d$ to deduce that

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

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

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

$$\int_{\mathbf{R}^d} g^d \, dV \leq c'_d \alpha^{d/(d-1)} N, \tag{2}$$

where c'_d depends on d.

4 Construction of a continuous separator

In this section we find a sphere S such that

$$\int_S g^{d-1} \, dA = O(N^{(d-1)/d}) \tag{3}$$

and such that at most βN nodes are inside S and at most βN are outside, where $\beta < 1$ depends on d.

This sphere can be constructed using a theorem from Miller and Thurston. In particular, they have the following result.

Theorem 2 *Let T be a set of N nodes embedded in $\mathbb{R}^d$, and let $g : \mathbb{R}^d \to \mathbb{R}$ be a nonnegative-valued measurable function with compact support. Suppose that*

$$\int_{\mathbf{R}^d} g^d \, dV = h. \tag{4}$$

Then there exists a sphere S such that at most βN nodes of T are strictly inside S, at most βN are strictly outside, and

$$\int_S g^{d-1} \, dA \leq \phi_d h^{(d-1)/d}, \tag{5}$$

where $\beta = (d+1)/(d+2)$ and ϕ_d is a positive constant depending on d.

The proof of their theorem uses conformal mappings and the notion of the "centerpoint" of a set of points (see below). It concludes with the use of the Hölder inequality.

In this section we propose two potentially useful variants to this this result. The first variant is simpler because the proof technique requires only counting (plus the Hölder inequality), although the constant factor β is much worse (exponentially close to 1).

The second variant is notable because it attains a better constant value ($\beta = d/(d+1)$) and because it is very reminiscent of the Lipton and Tarjan's [1979] planar separator theorem, one of the first general-purpose results on separators. The separator, however, has a more complicated shape than a sphere.

First, we prove the result with $\beta = 1-2^{-d-1}$ using a counting argument. Let $\hat{G}$ denote the set of N points in $\mathbb{R}^d$ at which nodes of G are embedded. We first construct a sequence of numbers $b_1, \ldots, b_d$ such that the set:

$$F_1 = \hat{G} \cap \{(x_1, \ldots, x_d) : x_1 \leq b_1, \ldots, x_d \leq b_d\}$$

has at least $N/2^d$ points, as does the set

$$F_2 = \hat{G} \cap \{(x_1, \ldots, x_d) : x_1 \geq b_1, \ldots, x_d \geq b_d\}.$$

In order to construct these two sets, rigid motions of $\mathbb{R}^d$ are permitted.

The numbers $b_1, \ldots, b_d$ are constructed inductively; suppose $b_1, \ldots, b_{i-1}$ have been constructed so that $\hat{G} \cap A$ and $\hat{G} \cap B$ each have at least $N/2^{i-1}$ points where

$$A = \{(x_1, \ldots, x_d) : x_1 \leq b_1, \ldots, x_{i-1} \leq b_{i-1}\}$$

and

$$B = \{(x_1, \ldots, x_d) : x_1 \geq b_1, \ldots, x_{i-1} \geq b_{i-1}\}.$$

(Note: for the inductive base case, we can take b_1 to be the median x_1 coordinate among points in $\hat{G}$. Accordingly, we can assume $i \geq 2$ in the following.) Then we take the median value of x_i coordinates among points in $\hat{G} \cap (A \cup B)$. Let this median value be b_i. Let L and U denote the halfspaces $\{x_i \leq b_i\}$ and $\{x_i \geq b_i\}$ respectively. Then one of the two following statements must be true:

1. There are at least $N/2^i$ points in $\hat{G} \cap A \cap L$ and at least $N/2^i$ in $\hat{G} \cap B \cap U$.

2. There are at least $N/2^i$ points in $\hat{G} \cap A \cap U$ and at least $N/2^i$ in $\hat{G} \cap B \cap L$.

In particular, suppose both statements were false, e.g., suppose that $\hat{G} \cap A \cap L$ and $\hat{G} \cap B \cap L$ each had fewer than $N/2^i$ points. This means that $\hat{G} \cap (A \cup B) \cap L$ has fewer than $N/2^{i-1}$ points, which contradicts the choice of b_i as the median. Other ways of negating both of the two statements also lead to contradictions.

Thus, either statement 1 or 2 holds. If statement 1 holds, this exactly proves the inductive statement. If statement 2 holds, then we reflect the embedding by transforming coordinate x_i of all nodes to $2b_i - x_i$. After this reflection, statement 2 also proves the inductive statement.

Therefore, the induction argument shows that the sequence $b_1, \ldots, b_d$ exists. Next, let ζ be the median distance of points in $\hat{G} \cap (F_1 \cup F_2)$ from the point $b = (b_1, \ldots, b_d)$. Using a similar reflection argument, we can conclude that there are at least $N/2^{d+1}$ points in each the two sets

$$G_1 = \hat{G} \cap \{(x_1, \ldots, x_d) : x_1 \leq b_1, \ldots, x_d \leq b_d, \|x - b\|_2 \leq \zeta\}$$

and

$$G_2 = \hat{G} \cap \{(x_1, \ldots, x_d) : x_1 \geq b_1, \ldots, x_d \geq b_d, \|x - b\|_2 \geq \zeta\}.$$

See Figure 4 for an illustration of the boundaries of these two sets in two dimensions. We claim that it is possible to surround G_1 with two concentric spheres S_1 and S_2 of radii $r_1 < r_2$ such that

1. Set G_1 is in the interior of S_1.

2. Set G_2 is in the exterior of S_2.

3. The ratio r_1/r_2 depends only on d.

The two spheres are illustrated in Figure 4. The center c of the two spheres may be chosen to be the point

$$(b_1 - \zeta/d, \ldots, b_d - \zeta/d).$$

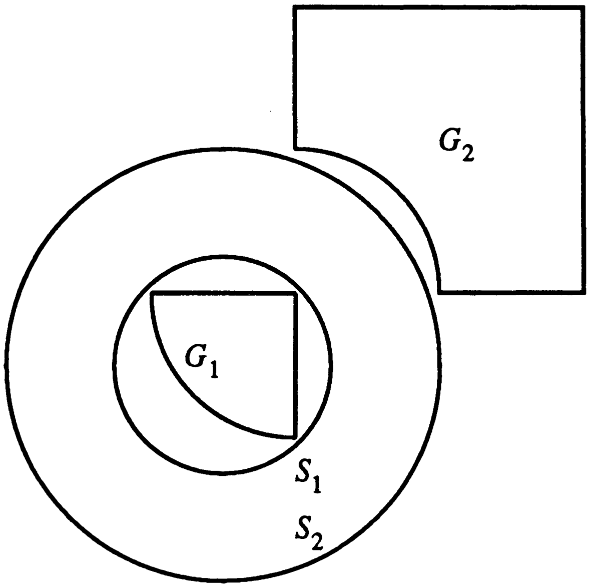

Figure 4: The construction from the counting argument in two dimensions.

Then one can compute that setting $r_1 = \zeta(1 - 1/d)$ and $r_2 = \zeta(1 + 1/d)$ satisfies all the conditions.

Finally, we construct the sphere to satisfy (5). The previous construction shows that the ratio r_1/r_2 is a positive constant $\gamma_d < 1$. Let S_r denote (the surface of) the sphere of radius r centered at c. Then we consider the expected value of

$$\int_{S_r} g^{d-1}\, dA$$

where r is chosen uniformly at random between r_1 and r_2. Let a be this average value, which is exactly equal to

$$a = \frac{1}{r_2 - r_1} \int_{r_1}^{r_2} \int_{S_r} g^{d-1}\, dA\, dr.$$

The double integral above can be interpreted as a volume integral of g^{d-1} over the region $P \in \mathbb{R}^3$ contained between S_1 and S_2, so we have

$$a = \frac{1}{r_2 - r_1} \int_P g^{d-1}\, dV.$$

Now we apply the Hölder inequality. The Hölder inequality says that for nonnegative functions ϕ and ψ suitably integrable on a measurable set P

and for positive real numbers p, q such that $1/p + 1/q = 1$, the following relation holds:

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

See, for example, Royden [1968]. We apply this function to our problem with P as above, $p = d/(d-1)$, $q = d$, $\phi = g^{d-1}$ and $\psi = 1$ (constant function) on P. Then we conclude that

$$a \leq \frac{1}{r_2 - r_1}\left(\int_P g^d\, dV\right)^{(d-1)/d} \cdot \left(\int_P 1\, dV\right)^{1/d}.$$

The first parenthesized integral is bounded by h as in (4). The second parenthesized integral is the volume of P. The volume of P is $v_d(r_2^d - r_1^d)$. An upper bound on the dth root of this volume is $v_d^{1/d} r_2$. Combining everything gives

$$a \leq \frac{v_d^{1/d} r_2}{r_2 - r_1} h^{(d-1)/d}.$$

The fraction in the previous expression is equal to a constant depending only on d.

This shows that the expected value of

$$\int_{S_r} g^{d-1}\, dA \tag{6}$$

is bounded by $c'h^{(d-1)/d}$ if r is chosen uniformly at random from the interval $[r_1, r_2]$.

This is not exactly what is needed to show the existence of an efficient probabilistic algorithm; instead, we should show that the probability that an unacceptably large separator is chosen is small. We claim that

$$\text{Prob}\left(\int_{S_r} g^{d-1}\, dA > 2c'h^{(d-1)/d} : r \in [r_1, r_2]\right) \leq 1/2.$$

In particular, since the expected (average) value of (6) is at most $c'h^{(d-1)/d}$ and the integral is never smaller than 0, then no more than $1/2$ of the probability distribution could lie beyond $2c'h^{(d-1)/d}$.

Therefore, if we choose k values of r lying in $[r_1, r_2]$ at random, then the probability all the integrals of the form (6) will have value more than

$2c'h^{(d-1)/d}$ is at most 2^{-k}, which can be made arbitrarily small. Finally, recall that h was defined to be the integral of g^d in $\mathbb{R}^d$, which we proved last section was $c\alpha^{(d/(d-1)}N$, with a constant depending on d.

Thus, the algorithm to construct a sphere S satisfying (5) with at most $(1-2^{-d-1})N$ nodes of G inside or outside of S is as follows. First, the point $(b_1, \ldots, b_d)$ and the radius ζ are computed using the inductive procedure described above. This requires $d+1$ median-finding computations (see Aho, Hopcroft and Ullman [1974] for median-finding algorithms). Each median operation requires $O(N)$ steps. From this data the two radii r_1 and r_2 are computed, and then a random number r is selected uniformly at random from the interval $[r_1, r_2]$. The sphere S_r could be tested directly to see whether it satisfies (5), but it is easier to use the procedure described in the next section to compute a node separator. The node separator can be tested to see if it is too large.

We now turn to the second argument to establish a variant of Theorem 2. This approach will yield a better constant at the expense of a more complicated separator shape. This approach, when viewed abstractly, is very reminiscent of Lipton and Tarjan's proof of the planar separator theorem. This approach is nonconstructive in that we do not have algorithms for every step.

First, we need to define the centerpoint of a finite set of nodes.

Definition 2 *Given a set $Q \subset \mathbb{R}^d$ of N nodes, a* **centerpoint** *of Q is a point $c \in \mathbb{R}^d$ such that for any halfspace H not containing c, there are at most $\frac{d}{d+1}N$ points in $Q \cap H$.*

Efficient algorithms for computing centerpoints are described in Miller and Teng [1990]. Assume that we have a centerpoint $c \in \mathbb{R}^d$ of the nodes of G embedded in $\mathbb{R}^d$.

Let $C(r)$ be the sphere of radius r around c. Define a function p of r to be the surface integral

$$p(r) = \int_{C(r)} g^{d-1}\, dA.$$

Now, pick r_1 and r_2 as follows. Let r_1 be the maximum radius such that

$$p(r_1) \leq \eta h^{(d-1)/d} \tag{7}$$

and such that at most $dN/(d+1)$ nodes are strictly interior to $C(r_1)$. Here, η is a constant that depends on d. The value for η is given below.

We note that this r_1 exists for the following reason. We can identify a particular r^* such that there are at most $dN/(d+1)$ nodes on strictly inside $C(r_1)$ if and only if $r \leq r^*$. Next, notice that $p(r)$ is a continuous function except possibly in the case that the centerpoint c coincides with an embedded node of G. (In this exceptional case we drop the term in the formula for g corresponding to this anomalous node; this does not affect the final separator bounds). Since p is continuous on $[0, r^*]$ and $p(0) = 0$, we conclude that r_1 exists.

Similarly, let r_2 be the minimum radius such that $p(r_2) \leq \eta h^{(d-1)/d}$ and such that at most $dN/(d+1)$ nodes are embedded strictly outside $C(r_2)$.

We now have some cases.

Case 1, $r_1 \geq r_2$. In this case, there are at least $N/(d+1)$ nodes on or outside $C(r_1)$ and at least $N/(d+1)$ on or inside $C(r_2)$. Therefore, sphere $C(r_1)$ (as well as $C(r_2)$) satisfies the theorem.

If we are not in Case 1, then $\mathbb{R}^d$ is partitioned into three regions A, B, C, namely, the inside of $C(r_1)$, the space between $C(r_1)$ and $C(r_2)$, and the outside of $C(r_2)$. Formally,

$$\begin{aligned} A &= \{x \in \mathbb{R}^d : \|x - c\| < r_1\}, \\ B &= \{x \in \mathbb{R}^d : \|x - c\| \in [r_1, r_2]\}, \\ C &= \{x \in \mathbb{R}^d : \|x - c\| > r_2\}. \end{aligned}$$

Because of the choices of r_1 and r_2, we know that at least $N/(d+1)$ nodes are in $B \cup C$ and at least $N/(d+1)$ are in $A \cup B$.

Case 2. At most $N/(d+1)$ nodes are in B. Then $A \cup C$ has at least $dN/(d+1)$, so at least one of A or C has at least $N/(d+1)$ (provided $d \geq 2$). Say, for example, that C has at least $N/(d+1)$. We already know $A \cup B$ has at least $N/(d+1)$. Therefore, $C(r_2)$ is a valid separator.

Case 3. B has between $N/(d+1)$ and $dN/(d+1)$ nodes. In this case our separator is $C(r_1) \cup C(r_2)$, and our two subdomains are B and $A \cup C$.

Case 4. B has more than $dN/(d+1)$ nodes. This is the long case. To simplify the notation for this case we assume that $c = 0$, i.e., the centerpoints of the nodes is at the origin.

First note that by choice of r_1, r_2, the integral of g^{d-1} on $C(r)$ for any $r \in [r_1, r_2]$ is at least $\eta h^{(d-1)/d}$. We now compute an upper bound on

$(r_2 - r_1)/r_1$. Observe that

$$\begin{aligned}
h^{(d-1)/d} v_d^{1/d} r_2 &\geq \left(\int_B g^d \, dV\right)^{(d-1)/d} \cdot \mathrm{vol}(B)^{1/d} \\
&\geq \int_B g^{d-1} \, dV \\
&\geq \int_{r_1}^{r_2} \int_{C(r)} g^{d-1} \, dV \, dr \\
&\geq (r_2 - r_1)\eta h^{(d-1)/d}.
\end{aligned}$$

Clearing the factor $h^{(d-1)/d}$ and assuming $\eta = 2v_d^{1/d}$ we conclude that $r_2 - r_1 \leq r_2/2$, i.e., $r_2 \leq 2r_1$. In the above chain of inequalities we used facts derived earlier plus the Hölder inequality.

Let U be the unit sphere in $\mathbb{R}^d$, that is, the set of vectors of unit length. For a nonzero vector a, let $a^\perp$ be the hyperplane through the origin normal to a. Finally, let rU denote the vectors of length r.

Suppose a chosen uniformly at random from U; then we have the following chain of equations about the expected value of the integral of g^{d-1} on $a^\perp \cap B$ (note that $a^\perp \cap B$ is the intersection of a hyperplane through the origin with the volume between two spheres centered at the origin). In this chain, s_d denotes the surface area of U, and τ denotes the expected value.

$$\begin{aligned}
\tau &= E\left(\int_{x \in a^\perp \cap B} g(x)^{d-1} \, dA\right) \\
&= \frac{1}{s_d} \int_{a \in U} \int_{x \in a^\perp \cap B} g(x)^{d-1} \, dA \, dA' \\
&= \frac{1}{s_d} \int_{a \in U} \int_{r_1}^{r_2} \int_{x \in a^\perp \cap rU} g(x)^{d-1} \, da \, dr \, dA' \\
&= \frac{1}{s_d} \int_{r_1}^{r_2} \int_{a \in U} \int_{x \in a^\perp \cap rU} g(x)^{d-1} \, da \, dA' \, dr \\
&= \frac{1}{s_d} \int_{r_1}^{r_2} r^{d-2} \int_{a \in U} \int_{w \in a^\perp \cap U} g(rw)^{d-1} \, da \, dA' \, dr && (8) \\
&= \frac{1}{s_d} \int_{r_1}^{r_2} r^{d-2} \int_{w \in U} \int_{a \in w^\perp \cap U} g(rw)^{d-1} \, da \, dA' \, dr && (9)
\end{aligned}$$

The interchange of a and w between (8) and (9) is explained in an appendix to this section. Notice that the integrand $g(rw)^{d-1}$ does not depend on the innermost variable of integration in (9). This means that we can replace that integration with the volume of $w^\perp \cap U$, which is s_{d-1}.

$$
\begin{aligned}
\tau &= \frac{s_{d-1}}{s_d}\int_{r_1}^{r_2} r^{d-2}\int_{w\in U} g(rw)^{d-1}\,dA\,dr \\
&= \frac{s_{d-1}}{s_d}\int_{r_1}^{r_2} r^{-1}\int_{x\in rU} g(x)^{d-1}\,dA\,dr \\
&= \frac{s_{d-1}}{s_d}\int_{x\in B} \|x\|^{-1} g(x)^{d-1}\,dV \\
&\le \frac{s_{d-1}}{s_d r_1}\int_{x\in B} g(x)^{d-1}\,dV \\
&\le \frac{s_{d-1}}{s_d r_1}\left(\int_{x\in B} g(x)^d\,dV\right)^{(d-1)/d}\cdot\left(\int_{x\in B} 1\,dV\right)^{1/d} \\
&\le \frac{s_{d-1}}{s_d r_1} h^{(d-1)/d}\,\mathrm{vol}(B)^{1/d} \\
&\le \frac{s_{d-1}}{s_d r_1} h^{(d-1)/d}(v_d {r_2}^d)^{1/d} \\
&\le \frac{v_d^{1/d} s_{d-1}}{s_d r_1} h^{(d-1)/d} r_2 \\
&\le \frac{2 v_d^{1/d} s_{d-1}}{s_d} h^{(d-1)/d}.
\end{aligned}
$$

In the above chain of inequalities we used the Hölder inequality, the bound on the integral of g^d on $\mathbb{R}^d$, the fact that B is contained in a ball of radius r_2, and the fact that $r_2 \le 2r_1$.

Thus, if a is chosen at random, then with high probability the bound of the integral on $a^\perp \cap B$ is at most a constant times $h^{(d-1)/d}$. Accordingly, assume that a has been chosen with this property.

Since c is a centerpoint, at most $(d-1)N/d$ of the nodes are on either side of $a^\perp$. Let the two sides of this hyperplane be denoted by E and F; they form a partition of $\mathbb{R}^d$ that is "orthogonal" to the partition given by A, B, C.

Now, we take cases. The first case is that $B \cap E$ has fewer than $N/(d+1)$ nodes. Then one of our subdomains is $B \cap F$ and the other subdomain is the complement of $B \cap F$. Notice that since B has at least $dN/(d+1)$ nodes and $B \cap E$ has fewer than $N/(d+1)$, it must be the case that $B \cap F$ has at least $(d-1)N/(d+1)$, i.e., at least $N/(d+1)$. Moreover, the complement of $B \cap F$ contains E, which also has at least $N/(d+1)$. Thus, we have a partition satisfying the bound. The boundary of $B \cap F$ is made up of pieces of $C(r_1)$, $C(r_2)$, and $a^\perp \cap B$. This means that the integral of g^{d-1} on the boundary is bounded by a constant multiplied by $h^{(d-1)/d}$.

The second case is that $B \cap E$ has between $N/(d+1)$ and $dN/(d+1)$ nodes. Then our two subdomains are $B \cap E$ and $\mathbb{R}^d - (B \cap E)$. Clearly this partition has all the desired properties.

The final case is that $B \cap E$ has more than $dN/(d+1)$ nodes. This is impossible because at most $dN/(d+1)$ nodes can be on either side of $a^\perp$.

Appendix: Interchanging order of integrations.

In this appendix we explain the interchange of integrals between (8) and (9). We remark that the same interchange of integrals is required to establish the Miller-Thurston theorem and is used implicitly in that paper.

Let E be the following set:

$$E = \{(a, w) \in \mathbb{R}^{d+1} \times \mathbb{R}^{d+1} : \|a\| = \|w\| = 1; \langle a, w\rangle = 0\}.$$

Here, $\langle\cdot,\cdot\rangle$ denotes inner product. Thus, E is a subset of $U \times U$. For the rest of this appendix we denote the standard embedding of the d-dimensional sphere in $\mathbb{R}^{d+1}$ by S^d rather than U.

Observe that the two inner integrals in both (8) and (9) are both parametrizations of E; in other words

$$E = \bigcup_{a \in S^d} \left(\{a\} \times (a^\perp \cap S^d)\right) = \bigcup_{w \in S^d} \left((w^\perp \cap S^d) \times \{w\}\right).$$

Thus, the sets of integration are the same; all that remains is to compare the volume elements. Note that E has a natural volume element arising from its embedding in $\mathbb{R}^{2d+2}$. Instead of connecting the volume elements in (8) and (9) to each other, we compare them both to the natural volume element on E.

Analyzing the volume elements requires some differential geometry. Let π_1 defined on E be the projection onto the first $d+1$ coordinates. Define π_2 as the projection onto the last $d+1$ coordinates. Observe that the image of π_1 is S^d. Moreover, for any $a \in S^d$, $\pi_1^{-1}(a)$ is a copy of S^{d-1} embedded in $\mathbb{R}^{d+1}$. Under these circumstances, we say that (E, π_1) is a *bundle* whose *base* B is S^d (i.e., $B = S^d = \pi_1(E)$) and whose *fibers* are isometric to S^{d-1}. We denote the fiber of a particular $a \in B$ by $\mathrm{Fib}(a)$, i.e., $\mathrm{Fib}(a) = \pi_1^{-1}(a) = a^\perp \cap S^d$. Note that a small neighborhood of a point in E is homeomorphic to the neighborhood of a point in $S^d \times S^{d-1}$, so (E, π_1) is a "local product."

The integration in (8) corresponds to integrating on a "fiber-by-fiber" basis. If E were *trivial* bundle, i.e., E were diffeomorphic to $S^d \times S^{d-1}$, then we could convert the fiber integration to an integration over E simply by exhibiting the diffeomorphism as a function in terms of coordinates, and then finding the determinant of its Jacobian. Since E is nontrivial, however, a representation in terms of coordinates does not exist.

Instead, we produce a volume element by examining the tangent spaces. For a manifold M, let $T_x(M)$ denote the *tangent space* of M at point $x \in M$. The tangent space has formal definition in differential geometry. Since all of our manifolds come with a natural embedding in $\mathbb{R}^n$ for some n, we can use the intuitive definition that $T_x(M)$ for M embedded in $\mathbb{R}^n$ is the subspace of vectors in $\mathbb{R}^n$ tangent to M at x. If our manifold is given by constraints (as is E, for example), then the tangent space at point is the subspace of vectors orthogonal to the gradient of each constraint. By writing out these gradients, we can in fact deduce immediately that $T_{(a,w)}(E)$ is the $2d-1$-dimensional subspace of $\mathbb{R}^{2d+2}$ that is a direct sum of three mutually orthogonal vector spaces:

$$T_{(a,w)}(E) = T_1 \oplus T_2 \oplus T_3$$

where

$$\begin{aligned} T_1 &= \{(0,y) : \langle y, a\rangle = \langle y, w\rangle = 0\} \\ T_2 &= \{(x,0) : \langle x, a\rangle = \langle x, w\rangle = 0\} \\ T_3 &= \mathrm{span}((w, -a)). \end{aligned}$$

In general, if (E, π) is a bundle with base B, then there is a short exact sequence of tangent spaces arising from the natural injection and projection of tangent spaces at a point $(a, w) \in E$:

$$0 \longrightarrow T_w(\mathrm{Fib}(a)) \longrightarrow T_{(a,w)}(E) \longrightarrow T_a(B) \longrightarrow 0.$$

This short exact sequence gives rise to a linear isomorphism (not uniquely defined) given by

$$g : T_w(\mathrm{Fib}(a)) \times T_a(B) \to T_{(a,w)}(E).$$

The determinant of this isomorphism is uniquely defined (although the isomorphism itself is not). This determinant is plays the role of the determinant of the Jacobian in converting from integration over the fibers to integration over E.

For our manifolds, we can write out an explicit representation of the isomorphism. In particular,

$$T_w(\mathrm{Fib}(a)) = T_w(a^\perp \cap S^d) = \{y \in \mathbb{R}^{d+1} : \langle y, a\rangle = \langle y, w\rangle = 0\}$$

and for $w \in \mathrm{Fib}(a)$,

$$\begin{aligned} T_a(S^d) &= \{x \in \mathbb{R}^{d+1} : \langle x, a\rangle = 0\} \\ &= \{x \in \mathbb{R}^{d+1} : \langle x, a\rangle = \langle x, w\rangle = 0\} \oplus \mathrm{span}(w). \end{aligned}$$

There is a natural map g' from $T_w(\mathrm{Fib}(a))$ to T_1 above, and a map g'' from $T_a(S^d)$ to $T_2 \oplus T_3$. These maps respect the structure of the exact sequence. Combining these maps gives g. In particular, given (x, y) in the domain of g, we would decompose y as a sum $y' + y''$ where y' lies in $w^\perp$ and $y'' = cw$ for some scalar c, and then $g(x, y) = (0, y') + (x, 0) + c(w, -a)$. Since all these decompositions are orthogonal, it is easy to compute the determinant of g. It turns out that the determinant is $\sqrt{2}$ independent of a or w. This is because everything is isometrically mapped by g, except unit vector w is mapped to $(w, -a)$, a vector of length $\sqrt{2}$.

Thus, we see that the volume change in going from an integral over the fibers as in (8) to an integral over E is just a multiplication by $\sqrt{2}$. But now we observe that E has a second bundle structure, namely, the structure given by (E, π_2). Everything is symmetric, so integrating over E compared to the fibers in this bundle structure also introduces a factor of $\sqrt{2}$. This second bundle structure corresponds to the integral in (9). We have therefore proved that for an arbitrary integrable function f defined on E,

$$\sqrt{2}\int_{a\in S^d}\int_{w\in \mathrm{Fib}(a)} f(a, w) = \int_{(a,w)\in E} f(a, w)$$

and

$$\sqrt{2}\int_{w\in S^d}\int_{a\in \mathrm{Fib}(w)} f(a, w) = \int_{(a,w)\in E} f(a, w).$$

Therefore, the two integrals on the left hand sides above are equal.

5 Finding a node separator

In the last section we described three methods to find a surface such that no more than βN nodes are on either side of the surface, where β depends on d and is less than 1, and such that the integral of g^{d-1} on the surface is no more than a constant multiple of the integral of g^d in $\mathbb{R}^d$. For the function g constructed in Section 3, this separating surface S will satisfy

$$\int_S g^{d-1}\,dA \le c_d \alpha N^{(d-1)/d}$$

where c_d is a (new) constant depending on d.

In two of those methods, the separating set was a sphere; for this section we assume that we are in the sphere case.

Let r be the radius and c the center of the sphere S that satisfies (3). From this sphere we intend to find a set $G_\sim$ of nodes such that $G_1, G_2, G_\sim$ is a partition of G satisfying Theorem 1.

First, we give the rule for membership in $G_\sim$. Suppose sphere S crosses through an edge (i, j) of the graph, such that that $\pi(i)$ is inside S and $\pi(j)$ outside; let x be the intersection of the segment $(\pi(i), \pi(j))$ with S. If $\pi(i)$ is closer to x then we put i in $G_\sim$, else we put j in $G_\sim$. Also, if a node by coincidence is embedded exactly on S, then the node is put in $G_\sim$. The remaining nodes are put into either G_1 or G_2 according to the rule: nodes embedded inside S not in $G_\sim$ are put into G_1, and nodes outside S not in $G_\sim$ are put into G_2. It is clear that there is no edge from a node of G_1 to a node of G_2.

Also, by construction, neither G_1 nor G_2 has more than βN nodes. We now turn to the problem of establishing a bound on the size of $G_\sim$.

We first establish the following claim: if a node i is selected to be in $G_\sim$, then $\|\pi(i) - c\| \le 3r$. One possibility is that i is interior to S, in which case the distance between $\pi(i)$ and c is no more than r. The other possibility is that it is adjacent to a vertex i' embedded interior to S; since i was selected for $G_\sim$ rather than i', this implies that the distance from $\pi(i)$ to c is at most $3r$.

We next put an upper bound on the number of nodes of $G_\sim$ such that $\rho_i \ge r$ (recall that ρ_i is the distance from vertex $\pi(i)$ to its most distant graph neighbor). Suppose there are m such nodes; let M be the set of these nodes. Let i be such a node. Then we know that $\pi(i)$ is within distance $3r$

of c. On the other hand, there is a ball of radius $\rho_i/(2\alpha)$ around $\pi(i)$ that is disjoint from all other such balls for $i \in M$ because of the density condition. Therefore, a similar property is true of the m balls of radius $r/(2\alpha)$ centered at $\pi(i)$ for $i \in M$ since we are assuming that $r \leq \rho_i$. All these balls lie in a sphere of radius $3r + r/(2\alpha)$. Therefore, a volume argument the same as the argument used in Section 3 shows that

$$m \leq (6\alpha + 1)^d.$$

Thus, there are at most a constant number of nodes of $G_\sim$ that satisfy $\rho_i \geq r$.

Now we turn to the case of nodes of $G_\sim$ that satisfy $\rho_i < r$. Call this set W, and choose $i \in W$. We want to put a constant lower bound on

$$\int_S f_i^{d-1}\,dA. \tag{10}$$

Recall that the function f_i is $1/\rho_i$ inside a ball of radius ρ_i centered at $\pi(i)$ and 0 elsewhere. We claim that S passes at a distance no more than $\rho_i/2$ from $\pi(i)$; this follows from the inclusion of i in $G_\sim$. We consider the "patch" of S that is contained in the ball $\{x : \|x - \pi(i)\| \leq \rho_i\}$. We can show that because the radius of S is greater than ρ_i and the distance of the patch to the ball's center is no more than $\rho_i/2$, we can get a lower bound of $v_{d-1}(\sqrt{7}\rho_i/4)^{d-1}$ on the area of this patch.

Since f_i^{d-1} is $1/\rho_i^{d-1}$ everywhere on this patch, we conclude that integral (10) is at least $v_{d-1}(\sqrt{7}/4)^{d-1}$, which we will call γ_d for simplicity.

This is true for every $i \in W$. Therefore, we get an upper bound on the size of $G_\sim$ as follows:

$$|G_\sim| \leq T_1 + T_2$$

where

$$T_1 = (6\alpha + 1)^d$$

and

$$T_2 = \sum_{i \in W} \left(\frac{1}{\gamma_d} \int_S f_i^{d-1}\,dA \right).$$

Here, T_1 accounts for nodes that have $\rho_i \geq r$ and T_2 accounts for nodes that have $\rho_i < r$. Now we derive an upper bound on T_2 as follows.

$$T_2 \;=\; \int_S \frac{1}{\gamma_d} \sum_{i \in W} f_i^{d-1}\,dA$$

$$\leq \int_S \frac{1}{\gamma_d} \sum_{i=1}^{N} f_i^{d-1}\, dA$$

$$\leq \int_S \frac{1}{\gamma_d} g^{d-1}\, dA$$

We discussed the integral on the last line at the beginning of the section. This quantity has $(c_d/\gamma_d)\alpha N^{(d-1)/d}$ as an upper bound.

Thus, the number of nodes in $G_\sim$ returned by this construction is at most a constant plus an $c'_d \alpha N^{(d-1)/d}$ term, which concludes the proof of the main theorem.

6 Practical issues

Let us call the algorithm derived in this paper **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 $1/2$ the nodes, i.e., a ratio of $1 + o(1)$. There is a standard technique originally due to Lipton and Tarjan [1979] 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 in the bound on the size of the separator. Vavasis [1990] discusses this technique.

Another practical issue is a splitting into more than one subdomain. If the number of domains desired is a power of 2, this accomplished by applying **Strong-Split** recursively to get domains of the desired size. This approach can be generalized for a number of domains not a power of 2.

7 Partitioning a domain without a mesh

A problem of interest is, given a domain $\Omega \subset \mathbb{R}^d$ (without reference to node points) partition this domain. We assume that we have a function g defined on Ω that specifies the density of nodes once a division into elements is performed. (In particular, $g(x)$ is equal to $1/h$, where h is approximately the diameter of the element that will contain point $x \in \Omega$.)

This problem arises in the following context. In some cases one wishes to apply domain decomposition in which the mesh is tailored to the subdomain boundaries (see, for example, Bramble, Pasciak and Schatz). This is possible only if the subdomain boundaries are selected before the mesh is laid out.

Our methods apply to this problem of handling a continuous density function; in fact, we have reduced the discrete density problem to the continuous problem, solved the continuous problem, and the pulled the continuous solution back to the discrete regime. In Section 4 we would replace counting arguments concerning nodes by arguments that instead consider the integral of g^d over volumes.

The only difficulty for the continuous case is the need to be able to integrate powers of the density function g over various volumes and surfaces. If g is specified in a way that these calculations are tractable, then all our techniques go through.

8 Conclusions and open questions

It would be of interest to prove the existence of a simple algorithm to achieve a 50-50 split. As mentioned in Section 6, a 50-50 split can be achieved by applying Lipton and Tarjan's technique to our algorithm, but only at the expense of a large constant in the separator bound. It is also of interest to remove the randomness from the algorithm at hand without making it too much more complicated.

Our approach seems to have the best possible dependence on α. Vavasis's earlier work proves a lower bound of $\alpha N^{2/3}$ in the three-dimensional case, and this bound probably extends to other dimensions.

Finally, it is of interest to come up with a class of graphs embedded in three dimensions with bounded separator sizes characterized by topological properties instead of geometric properties. Lipton and Tarjan's work on planar separators was based entirely on combinatorial topology.

We remark that Miller, Teng and Vavasis have come up with a class of graphs that includes both density graphs and planar graphs as special cases. This work is in progress.

Acknowledgement

Thanks to Allen Back of Cornell and Greg Kuperberg of Berkeley for explaining many of the steps needed to obtain (9). Thanks to Bruce Hendrickson for suggesting how to improve our original definition of of g in order to get optimal dependence on α.

References

A. V. Aho, J. E. Hopcroft, and J. D. Ullman [1974], *The Design and Analysis of Computer Algorithms*, Addison Wesley, Reading.

M. J. Berger and S. Bokhari [1987], A partitioning strategy for nonuniform problems on multiprocessors, *IEEE Trans. Comp.*, C-36:570-580.

M. Bern, D. Eppstein and J. Gilbert [1990], Provably good mesh generation, *Proc. 31st Symp. on Foundations of Computer Science*, 231–241.

P. E. Bjørstad and O. B. Widlund [1986], Iterative methods for the solution of elliptic problems on regions partitioned into substructures, *SIAM J. Numer. Anal.*, 23:1097-1120.

J. H. Bramble, J. E. Pasciak, and A. H. Schatz [1986], An iterative method for elliptic problems on regions partitioned into substructures, *Math. Comp.* 46:361-9.

T. F. Chan and D. C. Resasco [1987], A framework for the analysis and construction of domain decomposition preconditioners, UCLA-CAM-87-09.

L. P. Chew [1989], Guaranteed Quality Triangular Meshes, Cornell C.S. TR 89-893.

J. A. George [1973], Nested dissection of a regular finite element mesh, *SIAM J. Numer. Anal.*, 10:345-363.

Alan George and J. W. H. Liu [1978], An automatic nested dissection algorithm for irregular finite element problems, *SIAM J. on Numer. Anal.* 15:1053-1069.

R. J. Lipton, D. J. Rose and R. E. Tarjan [1979], Generalized nested dissection, *SIAM J. Numer. Anal.*, 16:346-358.

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

G. Miller and S.-H.Teng [1990], Centerpoints and point divisions, manuscript.

G. L. Miller and W. Thurston [1990], Separators in two and three dimensions, *Proceedings of the 1990 Symp. on Theory of Computing*, ???-???.

V. Pan and J. H. Reif [1985a], Efficent parallel solution of linear systems, *Proc. 17th Symposium on Theory of Computing,* 143–152.

V. Pan and J. H. Reif [1985b], Extension of parallel nested dissection algorithm to path algebra problems, Technical Report TR-85-9, Computer Science Department, State University of New York at Albany.

A. Pothen, H. D. Simon, and K.-P. Liu [1989], Partitioning sparse matrices with eigenvectors of graphs, NASA Ames Research Center RNR-89-009.

H. L. Royden [1968], *Real Analysis, 2nd Edition*, Macmillan, New York.

S. A. Vavasis [1990], Automatic domain partitioning in three dimensions, Cornell C.S. TR 90-1082.