Density graphs and separators

Gary L. Miller*
Stephen A. Vavasis#

TO APPEAR IN THE SECOND ANNUAL ACM-SIAM SYMPOSIUM ON DISCRETE ALGORITHMS

Abstract

We propose a class of graphs that would occur naturally in three-dimensional finite-element problems, and we prove an $O(N^{2/3})$ bound on separators for this class of graphs. We also propose a simple randomized algorithm to find this separator in $O(N)$ time. Such an algorithm could 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 interested in finding a map $u : \Omega \to \mathbb{R}$ such that

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

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 partition the region into subdomains. This is the goal of the paper at hand. The partitioning is accomplished

*School of Computer Science, Carnegie Mellon University and Department of Computer Science, University of Southern California.

#Department of Computer Science, Cornell University.

†The first author's work is supported in part by National Science Foundation grant DCR-8713489. The second author's work is supported in part by the U.S. Army Research Office through the Mathematical Sciences Institute of Cornell University and by an N.S.F. Presidential Young Investigator award.

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

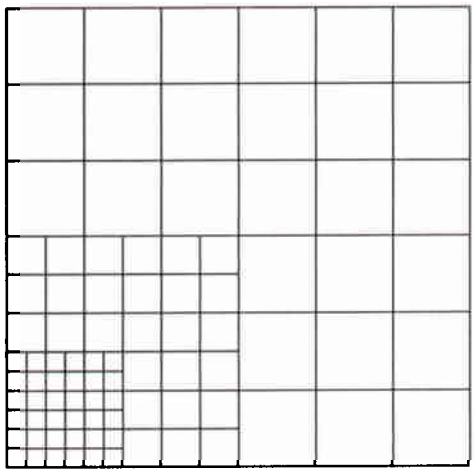

Figure 1: Berger and Bokhari's example of a density graph.

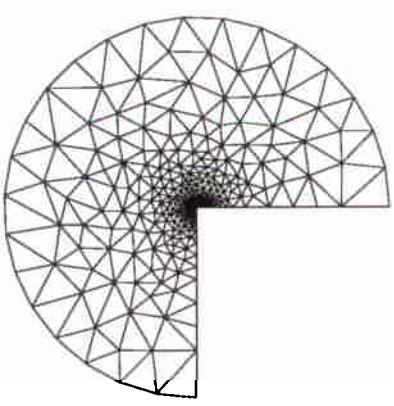

Figure 2: A graph generated by Chew's mesh generator.

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 two-dimensional analog 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 $O(N^{2/3})$ 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 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 $O(N^{2/3})$ (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 ran-*

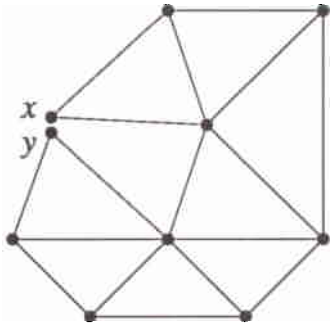

Figure 3: A stable graph with two nearby nodes.

domized $O(N)$ 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 $cN^{(d-1)/d}$ nodes, where 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 two 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 d_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:

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

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:

$$f(x) = \max(f_1(x), \ldots, f_N(x))$$

and

$$g(x) = \sum_{i=1}^{N} f_i(x).$$

We notice immediately that

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

because this integral is bounded by the sum of the integrals of the f_i^d.

We would like to establish a similar bound on the integral of g^d. This is the purpose of the following lemma.

Lemma 1 *For all x, the following inequalities hold:*

$$f(x) \le g(x) \le 2(4\alpha + 1)^d f(x).$$

Proof. The first inequality follows immediately from the definitions of f and g. 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, let i^* be the index such that $f_i(x)$ is maximum over all choices of i (break ties arbitrarily). Let t be the value of $1/d_{i^*}$. Then we ask, for integer $j \ge 0$, how many of the f_i can satisfy

$$2^{-j-1}t < f_i(x) \le 2^{-j}t? \tag{2}$$

Notice that all i such that $f_i(x) > 0$ will satisfy exactly one inequality of form (2). Suppose there are m_j distinct values, say set M_j, of indices i satisfying (2) for a particular choice of j. Then $\pi(i)$ for $i \in M_j$ must lie within distance $2^{j+1}/t$ of x by definition of f_i. On the other hand, the longest edge adjacent to $\pi(i)$ is of length at least $2^j/t$, also by definition of f_i. Surround each point i for $i \in M_j$ with a ball of radius $2^j/(2t\alpha)$. These balls must be disjoint for the following reason. The points $\pi(i)$ for $i \in M_j$ must be distance at least $2^j/(t\alpha)$ from one another by definition of a density graph. Accordingly, there are m_j disjoint balls of radius at least $2^j/(2t\alpha)$ around these m_j points. All of these balls lie in a sphere of radius $2^{j+1}/t + 2^j/(2t\alpha)$, that is, $2^{j+1}(1 + 1/(4\alpha))/t$ around x since the centers of the balls are within $2^{j+1}/t$.

The volume of each ball is $[v_d/(4\alpha)^d] \cdot (2^{j+1}/t)^d$, and the volume of the enclosing sphere is at most $v_d(1 + 1/(4\alpha))^d(2^{j+1}/t)^d$. This gives an upper

bound on m_j because the balls are disjoint. In particular, m_j is bounded by the quotient of these quantities, i.e.

$$\begin{aligned} m_j &\leq \frac{(1+1/(4\alpha))^d}{1/(4\alpha)^d} \\ &\leq (4\alpha+1)^d. \end{aligned}$$

Therefore, there are at most a constant number of values of i such that (2) is satisfied for any choice of j. The contribution to $f(x)$ from a particular i for a particular choice of a j is at most $2^{-j}t$. Therefore,

$$\begin{aligned} g(x) &\leq \sum_{j=0}^{\infty}(4\alpha+1)^d 2^{-j}t \\ &\leq 2t\cdot(4\alpha+1)^d. \end{aligned}$$

Since t was chosen to be equal to $f(x)$, the lemma follows. ∎

Therefore, g^d is no more than a constant multiple of f^d, where the constant is determined by the previous lemma. By (1) we have an equation of the form

$$\int_{\mathbf{R}^d} g^d\, dV = cN \tag{3}$$

where $c < q(\alpha, d)$ for some function q.

4 Construction of a separator

The next step in the proof of Theorem 1 is to find a $d-1$ dimensional sphere S embedded in $\mathbb{R}^d$ such that

$$\int_S g^{d-1}\, dA = O(N^{(d-1)/d}) \tag{4}$$

and such that at most βN nodes are inside S and at most βN are outside for some $\beta \in (0,1)$ that depends on d. For this we rely on a theorem of Miller and Thurston telling us that such a sphere may be found with constant $\beta = (d+1)/(d+2)$. The theorem is proved using conformal mapping and node centers. A simple counting argument gives the result with $\beta = 1 - 2^{-d-1}$. Another counting argument gives the result with $\beta = d/(d+1)$, but for this result S could have a shape more complicated than a sphere. In all three cases the Hölder inequality is required to bound integrals of g^{d-1} in terms of integrals of g^d. Also, the running time is randomized linear time (and also randomized NC) for all three approaches.

We briefly review the Miller-Thurston approach. The nodes of G are conformally embedded on the unit sphere in $\mathbb{R}^{d+1}$ (call this sphere Σ). A great circle on Σ is selected uniformly at random; this great circle is a sphere S when mapped back to $\mathbb{R}^d$, and the expected value of integral (4) for sphere S is $cN^{(d-1)/d}$ where c depends on d. Since the integral is never smaller than 0, no more than 1/2 of the probability distribution for the integral lies beyond $2cN^{(d-1)/d}$. This means that the probability that the left-hand side of (4) exceeds $2cN^{(d-1)/d}$ on k randomized attempts is 2^{-k}.

Let r be the radius and c the center of the sphere S that satisfies (4). From this sphere we intend to find a 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 $G_\sim$. 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 $(d+1)N/(d+2)$ 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 $\|\pi(i) - c\| \leq 3r$. 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 first put an upper bound on the number of nodes of $G_\sim$ such that $d_i \geq r$ (recall that d_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 $d_i/(2\alpha)$ 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 $r \leq d_i$. 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 $d_i \geq r$.

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

$$\int_S f_i^{d-1}\, dA. \tag{5}$$

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

Since f_i^{d-1} is $1/d_i^{d-1}$ everywhere on this patch, we conclude that integral (5) 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 $d_i \geq r$ and T_2 accounts for nodes that have $d_i < r$. Now we derive an upper bound on T_2 as follows.

$$\begin{aligned} T_2 &= \int_S \left(\frac{1}{\gamma_d} \sum_{i \in W} f_i^{d-1} \right) dA \\ &\leq \int_S \frac{f_1^{d-1} + \cdots + f_N^{d-1}}{\gamma_d}\, dA \\ &\leq \int_S \frac{(f_1 + \cdots + f_N)^{d-1}}{\gamma_d}\, dA. \end{aligned}$$

The integrand on the last line is exactly equal to g^{d-1}/γ_d by definition of g. Therefore, as derived at the beginning of the section, this quantity has $2cN^{(d-1)/d}/\gamma_d$ as an upper bound.

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

5 Future work and open questions

In future work we will apply our separator theorems to problems in computational geometry. With S.-H. Teng we have extended density graphs to a class of graphs that encompasses planar graphs as well; this will be reported in upcoming work.

It would be of interest to prove the existence of a simple algorithm to achieve a 50-50 split. A 50-50 split can be achieved using a standard recursive technique due to Lipton and Tarjan [1979] on our separator 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.

In addition, our approach in this paper does not seem to have optimal dependence on α; the constant factor in the $O(N^{(d-1)/d})$ separator bound is proportional to α^{2d} if we trace through all the constants. In upcoming work we show how to tighten the arguments so that the separator size is at most $c_d \alpha N^{(d-1)/d} + c'_d \alpha^d$; the leading term in this bound is apparently optimal.

Finally, it is of interest to come up with a class of graphs embedded in $\mathbb{R}^d$ 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.

References

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. L. Miller and W. Thurston [1990], Separators in two and three dimensions, *Proc. 22nd Symp. on Theory of Computing*, 300–309.

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.

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