Lower Bounds for Graph Embeddings and Combinatorial Preconditioners

Gary L. Miller and Peter C. Richter*
Department of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
{glmiller,prichter}@cs.cmu.edu

ABSTRACT

Given a general graph G, a fundamental problem is to find a spanning tree H that best approximates G by some measure. Often this measure is some combination of the congestion and dilation of an embedding of G into H. One example is the *routing time* $\rho(G, H) \leq O(\text{congestion} + \text{dilation})$, the number of steps necessary to route pairwise demands G on network links H in the store-and-forward packet routing model. Another is the *condition number* $\kappa_f(G, H) \leq O(\text{congestion} \cdot \text{dilation})$, the square root of which bounds the number of iterations necessary to solve a linear system with coefficient matrix G preconditioned by H using the classical conjugate gradient method. The algorithmic applications of being able to find (efficiently) a good tree approximation H for a graph G are numerous; but what if no good tree exists?

In this paper, we seek to identify the class of graphs G which are intrinsically difficult to approximate by a particular measure. It is easily seen that with respect to routing time, G is hardest to approximate by a tree H precisely when it contains either long cycles (which yield high dilation) or large separators (which yield high congestion). We show that with respect to condition number, the existence of long cycles or large separators in G is sufficient but *not* necessary for it to be hardest to approximate, by demonstrating a nearly-linear lower bound for the case in which G is a square mesh. The proof uses concepts from circuit theory, linear algebra, and geometry, and it generalizes to the case in which H is a spanning subgraph of G of Euler characteristic k. The result has consequences for the design of preconditioners for symmetric M-matrices and perhaps also of communication networks.

*The authors are supported in part by NSF Grants CCR–9902091 and ACI–0086093.

SPAA'04, June 27–30, 2004, Barcelona, Spain.

Categories and Subject Descriptors

E.1 [**Data structures**]: Graphs and networks; F.2.2 [**Nonnumerical algorithms and problems**]: Routing and layout; G.1.3 [**Numerical linear algebra**]: Conditioning, linear systems (direct and iterative methods)

General Terms

Algorithms, Theory

Keywords

Congestion, dilation, graph embeddings, preconditioning

1. INTRODUCTION

The challenge of finding a tree H that best approximates a graph G is an important one in the design of efficient algorithms. The paradigm is straightforward: many computational problems admit algorithms that perform better on trees, due to their simple structure, than on general graphs. Solving such a problem on a tree approximation H rather than the original graph G is sometimes the best route to an exact or approximate solution of the problem on G. Problems of this sort arise in parallel algorithms (e.g., routing) and numerical algorithms (e.g., preconditioning).

1.1 Problems and Results

Let G be a general graph on n nodes; we wish to find a spanning tree H that best approximates G by some measure combining the congestion and dilation of an embedding of G into H. One such measure is the *routing time* $\rho(G, H) \leq O(\text{congestion} + \text{dilation})$, the number of steps necessary to route pairwise demands G on network links H in the store-and-forward packet routing model. Another is the *condition number* $\kappa_f(G, H) \leq O(\text{congestion} \cdot \text{dilation})$, the square root of which bounds the number of iterations necessary to solve a linear system with coefficient matrix G preconditioned by H using the classical conjugate gradient method.[1] Tree approximations are useful for routing on fixed-connection networks, iterative solution of linear systems, and many other computational problems, and efficient algorithms for finding them are known.

In this paper, we seek to identify the classes of graphs G which are *intrinsically* the most difficult to approximate

[1]We use the same notation to represent both a graph and the associated matrix, assumed to be the *Laplacian* or more generally an *M-matrix*, which we define shortly.

with respect to routing time or condition number; i.e., those for which every tree H is a relatively poor approximation of G.

It is easily seen that with respect to routing time, G is hardest to approximate by a tree H precisely when it contains either long cycles[2] (which yield high dilation) or large separators (which yield high congestion). For any graph G, there is a tree H such that $\rho(G,H) \leq O(n)$; if G is a cycle or an expander, this is optimal. If G is a square mesh, then the optimal H satisfies $\rho(G,H) = \Theta(\sqrt{n})$. These results can be generalized to the case in which H is a spanning subgraph of G of Euler characteristic k.[3] Bern, Gilbert, Hendrickson, Nguyen, & Toledo [2] showed (implicitly) that for any G, there is an H such that $\rho(G,H) \leq O(\frac{n}{\sqrt{k+1}})$, or $\rho(G,H) \leq O(\frac{n}{k+1})$ if G has bounded genus; we show that if G is a fractal with many long cycles or an expander, then $\rho(G,H) \geq \Omega((\frac{n}{k+1})^{1-o(1)})$. If G is a square mesh, then the optimal H satisfies $\rho(G,H) = \Theta((\frac{n}{k+1})^{1/2})$.

We show that with respect to condition number, the existence of long cycles or large separators in G is sufficient but *not* necessary for it to be hardest to approximate. Boman & Hendrickson [6] showed that for any G, there is a tree H such that $\kappa_f(G,H) \leq O(n^{1+o(1)})$; if G is a cycle or an expander, then every H satisfies $\kappa_f(G,H) \geq \Omega(n)$. We prove that if G is a square mesh, then every H satisfies $\kappa_f(G,H) \geq \Omega(n^{1-o(1)})$. The proof uses concepts from circuit theory, linear algebra, and geometry, and it generalizes to the case in which H is a spanning subgraph of G of Euler characteristic k. Spielman & Teng [19] showed that for any graph G, there is an H such that $\kappa_f(G,H) \leq O(\frac{n^{1+o(1)}}{\sqrt{k+1}})$, or $\kappa_f(G,H) \leq O(\frac{n^{1+o(1)}}{k+1})$ if G has bounded genus; we show that if G is a square mesh, then every H satisfies $\kappa_f(G,H) \geq \Omega((\frac{n}{k+1})^{1-o(1)})$. The same lower bound holds when G is a fractal with many long cycles or an expander.

Our results have consequences for the design of preconditioners for symmetric M-matrices, which arise from finite-difference and finite-element discretizations of elliptic partial differential equations, and perhaps also of communication networks.

1.2 Related Work

The problem of approximating a general network by a low-congestion tree, which has application to oblivious routing, has been studied by Räcke [17]. He gave a super-polynomial time algorithm to construct for any network a Steiner tree prescribing how to route demands obliviously so that the congestion is to within a polylogarithmic factor of optimal. Polynomial time algorithms were subsequently given by Bienkowski, Korzeniowski, & Räcke [3] and by Harrelson, Hildrum, & Rao [11]. The problem of approximating a finite metric space (i.e., weighted complete graph) by a low-distortion (i.e., low-dilation) tree, which has application to the design of approximation algorithms, has been studied by Rabinovich & Raz [16]. They showed that the metric represented by a simple cycle cannot be approximated to within sublinear distortion by a tree metric. Naidu & Ramesh [15] generalized this result to approximation by a metric of positive Euler characteristic. Gupta [9] showed that approximation by a tree metric is not asymptotically improved by the addition of Steiner nodes, leading to a simpler proof of the result of Rabinovich & Raz.

The problem of designing and analyzing combinatorial preconditioners for linear systems was pioneered by Vaidya [21]. Although he never published his results, they have since appeared and been extended in Joshi [12], Bern, Gilbert, Hendrickson, Nguyen, & Toledo [2], Boman & Hendrickson [5], and Boman, Chen, Hendrickson, & Toledo [4]. The development of Steiner tree preconditioners, which are typically more effective than spanning tree preconditioners, is due to Gremban, Miller, & Zagha [8], Gremban [7], and Maggs, Miller, Parekh, Ravi, & Woo [14]. Reif [18] was the first to give a nearly-linear time system solver based on combinatorial preconditioners, though the result was only shown to hold for certain bounded-degree, planar linear systems. Boman & Hendrickson [6] gave the best upper bound for preconditioning a general graph G with a spanning tree H, using the spanning trees of Alon, Karp, Peleg, & West [1]. Spielman & Teng [19] generalized their result to H with positive Euler characteristic. Spielman & Teng [20] developed a nearly-linear time system solver for certain linear systems using combinatorial preconditioners in combination with a recursive application of the conjugate gradient method.

1.3 Outline

The remainder of this paper is organized as follows. In Section 2, we introduce notation and concepts needed in the following sections. In Section 3, we illustrate two applications of our results, to packet routing and preconditioning. In Section 4, we present dilation and congestion lower bounds for fractals, expanders, and square meshes. In Section 5, we give nearly-linear condition number lower bounds for fractals, expanders, and square meshes. In Section 6, we offer some final remarks.

2. PRELIMINARIES

The following definitions and mathematical background on graphs, matrices, and electrical networks are prerequisite for understanding later sections.

2.1 Graphs and Matrices

Let $G = (V, E_G)$ be a loopless, undirected graph on n nodes. The *boundary* $\partial_G U$ of $U \subseteq V$ is the set of edges in E_G connecting U to $V \setminus U$. An *α-separator* is a set of edges $\partial_G U$ for some $U \subseteq V$ which satisfies $\min\{U, V \setminus U\} \geq \alpha n$. The *genus* of a graph is the minimum-genus surface onto which it can be mapped properly (i.e., without edge crossings). For example, a *planar* graph has genus zero. The *skeleton* of a graph is the subgraph that remains after all nodes of degree one or two have been eliminated (iteratively) via edge contraction.

We define a *vine* to be a set $U \subseteq V$ such that (i) $G|U$ (i.e., the restriction of G to U) is a spanning tree, and (ii) the number of nodes of U incident to $\partial_G U$ (the *fixed ends* of the vine) is either one or two. We say the vine is *fixed* if it has two fixed ends, and *loose* if it has only one. If the vine is loose, we can choose any leaf node of $G|U$ (except the fixed end) to be the *loose end*. The *stem* of the vine is the unique path in $G|U$ from one end of the vine to the other.

[2]Long *nearly isometric* cycles, to be precise. For instance, the perimeter of the planar square mesh is nearly isometric (to a simple cycle), but a Hamiltonian cycle of the mesh is not.

[3]The *Euler characteristic* of $H = (V, E_H)$ is given by $|E_H| - |V| + 1$.

A *fractal* is a graph which is self-similar (i.e., it is constructed from an initial graph by successively replacing edges in the graph by copies of the initial graph). An *expander* is a graph for which there is a constant $c > 0$ such that $|\partial_G U| \geq c \cdot \min\{|U|, |V \setminus U|\}$ for all $U \subseteq V$. The *planar square mesh* is the Cartesian product of a path with itself. The *toroidal square mesh* is the Cartesian product of a cycle with itself. When we refer to a *square mesh*, we mean either a planar or toroidal square mesh.

Let G be an $n \times n$ symmetric matrix. Its finite spectral *condition number* $\kappa_f(G)$ can be expressed by the following Courant-Fischer characterization:

$$\begin{aligned} \kappa_f(G) &= \max_{x \neq 0} \frac{x^T G x}{x^T x} \cdot \max_{x : Gx \neq 0} \frac{x^T x}{x^T G x} \\ &= (\max \lambda_f(G)) \cdot (\min \lambda_f(G))^{-1} \end{aligned} \quad (1)$$

Here $\max \lambda_f(G)$ and $\min \lambda_f(G)$ denote the maximum and minimum finite (nonzero) eigenvalues of G, respectively.

We say that G is an *M-matrix* if it is *diagonally dominant* (i.e., satisfies the property $G_{ii} \geq \sum_{j \neq i} |G_{ij}|$ for each i) and has nonpositive off-diagonal entries. In particular, an M-matrix is either nonsingular and positive definite, or it has a nullspace of the all-ones vector and is positive semidefinite. In the latter case, the M-matrix is said to be in *Laplacian* form.

Henceforth, any reference to a matrix form G of some graph G should be taken to mean the Laplacian form defined by $G_{ij} = -1$ if $(i,j) \in E_G$, $G_{ij} = 0$ if $(i,j) \notin E_G$, and $G_{ii} = \sum_{j \neq i} |G_{ij}|$.

2.2 Electrical Networks

Suppose we view two graphs $G = (V, E_G)$ and $H = (V, E_H)$ as (resistive) electrical networks, with nodes representing junctions or terminals and edges representing unit-size branch conductors. Let x be a vector defined on the nodes. Then $\mathcal{E}_G(x) = x^T G x = \sum_{(i,j) \in E_G} (x_i - x_j)^2$ and $\mathcal{E}_H(x) = x^T H x = \sum_{(i,j) \in E_H} (x_i - x_j)^2$ measure the power dissipation of the circuits G and H, respectively, when the potentials at the nodes are set according to the values of x. This is simply the power law: the power dissipation is equal to the conductance times the square of the potential difference. In keeping with the electrical analogy, we'll refer to the quadratic form $\mathcal{E}_G$ as the *energy burned* on G by the *potential vector* x. This gives some physical meaning to the finite spectral generalized *condition number* $\kappa_f(G, H)$, which is given by:

$$\begin{aligned} \kappa_f(G, H) &= \max_{x : Hx \neq 0} \frac{\mathcal{E}_G(x)}{\mathcal{E}_H(x)} \cdot \max_{x : Gx \neq 0} \frac{\mathcal{E}_H(x)}{\mathcal{E}_G(x)} \\ &= (\max \lambda_f(G, H)) \cdot (\min \lambda_f(G, H))^{-1} \end{aligned} \quad (2)$$

Here $\max \lambda_f(G, H)$ and $\min \lambda_f(G, H)$ denote the maximum and minimum finite (nonzero) generalized eigenvalues of the matrix G preconditioned by H.

When G has bounded degree and H is a spanning subgraph of H (as the case will be henceforth), it is easy to see that the eigenvalue $\min \lambda_f(G, H)$ must be $\Theta(1)$. In this case, to bound $\kappa_f(G, H)$ from below, it suffices to present a potential vector x on the nodes which burns many times more energy on G than on H. Of course, the potential vector x yielding the strongest possible lower bound is simply the eigenvector associated with the eigenvalue $\max \lambda_f(G, H)$; however, this eigenvector is typically difficult to determine analytically.

3. APPLICATIONS

As we have already suggested, the routing time and the condition number bound the running times of two important practical applications: routing on fixed-connection networks and iterative solution of linear systems.

3.1 Packet Routing

Consider the following packet routing problem. Let $G = (V, E_G)$ be a graph representing pairwise traffic demands; i.e., $(i,j) \in E_G$ if and only if nodes i and j wish to exchange packets. Let $H = (V, E_H)$ be a subgraph of G representing physical network links; i.e., $(i,j) \in E_H$ if and only if nodes i and j are connected by a pipe.[4] The task is to satisfy the demands G by routing the packets on the network H as quickly as possible. This is done using the *store-and-forward* model of packet routing: packet motion through the network is synchronous (at each step, each network link is traversed by either 0 or 1 packets), and packet storage is provided in the form of a constant-size queue residing at each link. A solution to the packet routing problem consists of two parts: the selection of a routing path for each packet, and a schedule for packet motion along each path. The first part is a mapping of each edge (i,j) of G to an (edge-simple) path in H between i and j, or an *embedding* $\varphi : G \hookrightarrow H$. Given an embedding φ and an optimal schedule for packet motion, we would like to know how many steps the routing solution takes.

The *dilation* $d_\varphi(G, H)$ of an embedding φ is the length of the longest path in H to which an edge of G is mapped. The *congestion* $c_\varphi(G, H)$ of φ is the largest number of such paths mapped through a particular edge of H, over all edges of H. To the network user, dilation and congestion represent a delay and a bottleneck, respectively; thus, we would like to minimize them. We call an H designed to minimize dilation or congestion a *spanner* or *filler*, respectively.

Both dilation and congestion are lower bounds on the number of steps taken by a routing solution. Perhaps surprisingly, they are also upper bounds: Leighton, Maggs, & Rao [13] proved a more general version of the following lemma:

LEMMA 3.1 (ROUTING LEMMA). *For each selection of routing paths $\varphi : G \hookrightarrow H$, there is an (optimal) schedule for packet motion that yields a solution to the packet routing problem in $\Theta(c_\varphi(G, H) + d_\varphi(G, H))$ steps.*

In particular, the *routing time* $\rho(G, H)$, the number of steps needed to solve the packet routing problem using an optimal selection of routing paths and an optimal schedule for packet motion, satisfies $\rho(G, H) = \Theta(\min_\varphi c_\varphi(G, H) + d_\varphi(G, H))$.

3.2 Preconditioning

Consider the problem of solving the linear system $Gx = b$, where G is the Laplacian matrix form of a graph G.[5] The

[4] If the edges of G missing from H represent down links, then the problem is one of fault tolerance. If H is restricted to be a tree, then the problem is one of finding the best subgraph on which routes are unique.

[5] Such a system arises, for example, when the Poisson equation $\Delta u = f$ with Neumann boundary conditions on a compact domain $\Omega \subset \mathbb{R}^d$ is discretized by finite differences, using the standard $(2d+1)$-point central difference scheme.

simplest solution technique is the *direct* method of Gaussian elimination: first compute the Cholesky factorization LL^T of G, then solve the lower- and upper-triangular systems $LL^T x = b$ by forward- and back-substitution. Naively, this takes time $O(n^3)$; typically, however, G has some special topological structure that permits faster solution. If G is a tree, then there is a *zero fill* elimination order taking time $O(n)$. If G is a square mesh, then there is a *nested dissection* elimination order taking time $O(n^{3/2})$. In the former case, Gaussian elimination is optimal. In the latter case, Gaussian elimination can be beaten by an *iterative* method, such as the conjugate gradient method.

The *conjugate gradient method* solves the system $Gx = b$ by minimizing the quadratic form $\phi(x) = \frac{1}{2}x^T Gx - x^T b$ over the orthogonal complement of the nullspace of G. It is a Krylov method: the ith approximation to x lies in the subspace spanned by the first i vectors of $\{b, Gb, \ldots, G^{n-1}b\}$. During each iteration of the method, the key operations (and costs) are a matrix-vector multiplication involving G and, typically, the direct solution of a linear system with coefficient matrix H. The matrix H, called the *preconditioner*, is a sparser approximation to G used to accelerate convergence. In the presence of roundoff, the preconditioned conjugate gradient method is guaranteed to converge to an ϵ-approximate solution (i.e., an x' satisfying $||Gx' - b|| \leq \epsilon||b||$) in $O(\sqrt{\kappa_f(G,H)}\log(\kappa_f(G)/\epsilon))$ iterations. The challenge of preconditioning is to find an H for which (i) Gaussian elimination on H is fast (faster than on G), and (ii) the condition number $\kappa_f(G,H)$ is small (smaller than $\kappa_f(G)$).

Vaidya [21] was the first to observe that by choosing H to be a good, relatively low-characteristic spanner and filler of G simultaneously, H can serve as a good *combinatorial preconditioner* of G.[6] Subsequently, Gremban [7] proved the *Congestion-Dilation Lemma*, a more general form of the following result:

LEMMA 3.2 (PRECONDITIONING LEMMA). *Let G be any graph and H be a spanning subgraph of G. Then $\kappa_f(G,H) \leq O(\min_\varphi c_\varphi(G,H) \cdot d_\varphi(G,H))$, where $\varphi : G \hookrightarrow H$.*

Unlike the routing lemma, the preconditioning lemma is not in general asymptotically tight. Bounding the condition number from below is typically more difficult than simply computing congestion or dilation, as we shall see.

4. LOWER BOUNDS FOR SPANNERS AND FILLERS

In this section, we prove lower bounds for low-characteristic spanners and fillers of sparse graphs, nearly matching the known upper bounds. We also introduce a scheme for partitioning meshes; the technique is critical for extending spanning tree results to spanning subgraphs of positive characteristic.

4.1 Fractals and Expanders

Intuitively, graphs with relatively poor spanners are that way because they contain long cycles. For example, the graph with the worst tree spanner is a simple cycle. As we show below, this principle generalizes: fractals with many long cycles have the worst spanners of positive characteristic. The fractals are based on those of Gupta, Krauthgamer, & Lee [10].

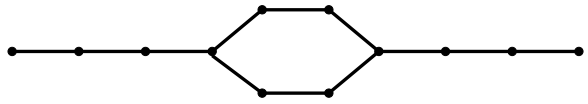

Figure 1: The graph F_1 in Theorem 4.1, for $i = 3$.

THEOREM 4.1. *There exists a bounded-degree planar fractal $F = (V, E_F)$ on n nodes of which every spanning subgraph $H = (V, E_H)$ of characteristic k yields $\min_\varphi d_\varphi(F,H) \geq \Omega((\frac{n}{k+1})^{1-o(1)})$.*

PROOF. The fractal F is constructed by the following recursive procedure. Let n and i be integers satisfying $n = (4r_i)^i - \sum_{j=1}^{i-1}(4r_i)^j$, where $r_i = \lceil(\frac{4}{3})^i\rceil$. Let F_1 be the graph consisting of four paths of length r_i glued together as shown in Figure 1, and for each $j \geq 2$ construct the graph F_j by taking a copy of F_{j-1} and replacing each edge with a copy of F_1. We define F to be the graph F_i. One can check that F is bounded-degree and planar and that it has n nodes and $(4r_i)^i$ edges.

Let $\mathcal{C}_j$ be the collection of all $(4r_i)^{j-1}$ (disjoint) cycles of F_j of length $2r_i$. Let $\mathcal{S}_j$ be the collection of all $(4r_i)^{j-1}$ (disjoint) connected subgraphs of F generated recursively from the $(4r_i)^{j-1}$ cycles of F_j in $\mathcal{C}_j$. In particular, each subgraph in $\mathcal{S}_j$ contains a simple cycle of length $2r_i(3r_i)^{i-j-1}$.

Let $t = \lceil\log_{4r_i}(k+1)\rceil$. At least one of the $(4r_i)^{t-1}$ subgraphs in $\mathcal{S}_t$, say S, must be such that $H|S$ is a tree. Let C be a simple cycle of length $2r_i(3r_i)^{i-t-1}$ in S with all but one edge (x,y) contained in H. Then the shortest path from x to y in H is the long way around C. Hence, the edge (x,y) suffers a dilation of $\frac{2}{3}(3r_i)^{i-t} - 1$, which is at least $(\frac{n}{k+1})^{1-1/\sqrt{\log_{4/3} n}}$. Thus, we have $\min_\varphi d_\varphi(F_i,H) \geq \Omega((\frac{n}{k+1})^{1-o(1)})$. □

Intuitively, graphs with relatively poor fillers are that way because they contain large separators. For example, the graph with the worst tree filler is an expander. As we show below, this principle generalizes: expanders have the worst fillers of positive characteristic as well. The result is likely folklore.

THEOREM 4.2. *Let $X = (V, E_X)$ be a bounded-degree expander on n nodes and $H = (V, E_H)$ be a spanning subgraph of X of characteristic k. Then $\min_\varphi c_\varphi(X,H) \geq \Omega(\frac{n}{k+1})$.*

PROOF. Let H' be a spanning tree of H. Since H' is bounded-degree, there is a single edge e of H' and a constant $\alpha > 0$ such that e is an α-separator of H'. Let this separator of H' partition the nodes into sets U and $V \setminus U$, say. Then the separators of H and X into U and $V \setminus U$ contain at most $k+1$ and at least $\Omega(n)$ edges, respectively. Thus, we have $\min_\varphi c_\varphi(X,H) \geq \Omega(\frac{n}{k+1})$. □

It follows immediately from Lemma 3.1 that every n-node fractal F and expander X satisfy $\rho(F,H) \geq \Omega((\frac{n}{k+1})^{1-o(1)})$ and $\rho(X,H) \geq \Omega(\frac{n}{k+1})$, respectively, for every spanning subgraph H of characteristic k.

[6] Moreover, iterative methods for solving linear systems using combinatorial preconditioners are highly parallelizable.

4.2 Square Meshes

The following lemma shows how to partition the nodes of a square mesh in such a way that spanning tree results can be extended to spanning subgraphs of positive characteristic.

LEMMA 4.3 (PARTITIONING LEMMA). *Let $M = (V, E_M)$ be a square mesh on n nodes and $H = (V, E_H)$ be a spanning subgraph of M of characteristic k. Then V contains a vine in H of size $\Omega(\frac{n}{k+1})$.*

PROOF. Obtain the skeleton H' of H by successively contracting edges to remove all nodes of degree one or two. Note that H' has $\Theta(k+1)$ nodes and edges and has characteristic k. Now color the nodes and edges of H' as follows. First, assign a unique color to each edge in H', and assign to each node in H' the color of one of its incident edges. Then, recover H by running the edge-contraction procedure used to derive H' from H in reverse. Each time an edge of H is recovered by subdividing an existing edge, assign the new node and edge the same color as the edge that was subdivided. Each time an edge of H is recovered by branching from an existing node, assign the new node and edge the same color as the node from which the branch occurred. Each of the $\Theta(k+1)$ color classes of nodes in H defines a vine. Hence, some vine has size $\Omega(\frac{n}{k+1})$. □

Lemma 4.3 enables us to prove dilation and congestion lower bounds for the square mesh.

THEOREM 4.4. *Let $M = (V, E_M)$ be a square mesh on n nodes and $H = (V, E_H)$ be a spanning subgraph of M of characteristic k. Then $\min_\varphi d_\varphi(M, H) \geq \Omega((\frac{n}{k+1})^{1/2})$ and $\min_\varphi c_\varphi(M, H) \geq \Omega((\frac{n}{k+1})^{1/2})$.*

PROOF. By Lemma 4.3, there exists a vine U in H of size $\Omega(\frac{n}{k+1})$. Since H has bounded degree and only two nodes of U are incident to $\partial_H U$, $|\partial_H U|$ is bounded. On the other hand, $|\partial_M U| \geq \Omega((\frac{n}{k+1})^{1/2})$, by the isoperimetric inequality in two dimensions. Hence, we have $\min_\varphi c_\varphi(M, H) \geq \Omega((\frac{n}{k+1})^{1/2})$. Moreover, since the diameter of U in M is $\Omega((\frac{n}{k+1})^{1/2})$, there is an edge of $\partial_G U$ at distance $\Omega((\frac{n}{k+1})^{1/2})$ from $\partial_H U$. It follows that $\min_\varphi d_\varphi(M, H) \geq \Omega((\frac{n}{k+1})^{1/2})$. □

Theorem 4.4 is tight: there exists a spanning subgraph H of characteristic k and an embedding $\varphi : M \hookrightarrow H$ for which $d_\varphi(M, H) \leq O((\frac{n}{k+1})^{1/2})$ and $c_\varphi(M, H) \leq O((\frac{n}{k+1})^{1/2})$, hence $\rho(M, H) \leq O((\frac{n}{k+1})^{1/2})$ by Lemma 3.1. The subgraph H is not unique; any characteristic k subgraph whose skeleton is a reasonably uniform coarser mesh and whose branches are reasonably good trees (e.g., those in Figures 2 and 3, the latter of which is due to Alon, Karp, Peleg, & West [1]) will be asymptotically optimal. The embedding φ mapping edges in M to shortest paths in H is sufficient here.

5. LOWER BOUNDS FOR COMBINATORIAL PRECONDITIONERS

In this section, we prove lower bounds for low-characteristic combinatorial preconditioners of sparse graphs, nearly matching the known upper bounds. Most importantly, we show that a simple square mesh is as hard as any graph to precondition.

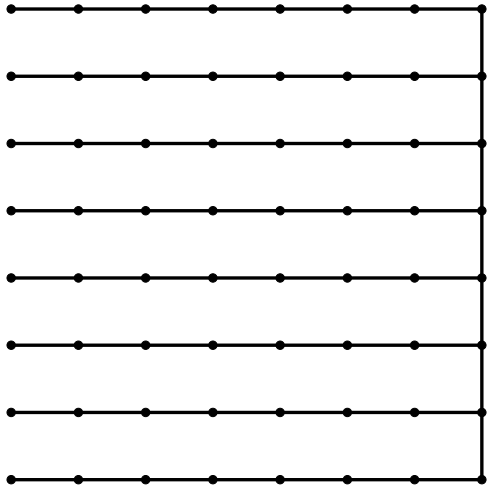

Figure 2: A good spanning tree for the square mesh.

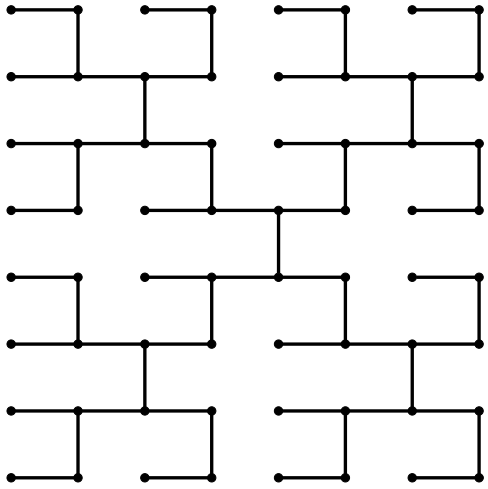

Figure 3: Another good tree for the square mesh.

5.1 Fractals and Expanders

Recall that in order to prove a lower bound on the condition number, we need to bound the energy burned by a graph (from below) and a spanning subgraph (from above) for a potential vector x.

For instance, let M be the square mesh and H be either of the spanning trees in Figures 2 and 3. For reference, let $\langle i, j\rangle$ be the node in the ith row and jth column. Then it is a simple puzzle to find a potential vector x which burns $\Omega(n)$ times more energy on M than on H. For the tree in Figure 2, one such x is the vector which sets the nodes on the path from $\langle 1, 1\rangle$ to $\langle 1, \sqrt{n}\rangle$ to potentials $\sqrt{n}-1, \ldots, 0$, respectively, and sets all other nodes to ground (zero potential). For the tree in Figure 3, one such x is the vector which sets the nodes on the path P_1 from $\langle \frac{1}{2}\sqrt{n}, \frac{1}{2}\sqrt{n}+1\rangle$ to $\langle \frac{1}{4}\sqrt{n}+1, \frac{1}{4}\sqrt{n}+1\rangle$ to potentials $0, \ldots, \sqrt{n}-1$, respectively, sets the nodes on the path P_2 from $\langle \frac{1}{2}\sqrt{n}+1, \frac{1}{2}\sqrt{n}+1\rangle$ to $\langle \frac{3}{4}\sqrt{n}, \frac{1}{4}\sqrt{n}+1\rangle$ to potentials $0, \ldots, \sqrt{n}-1$, respectively, and floats (or colors) the potential at each remaining node to the potential of the closest node in $P_1 \cup P_2$. Hence, $\kappa_f(M, H) \geq \Omega(n)$ if H is either of the trees in Figures 2

and 3. This generalizes to $\kappa_f(M,H) \geq \Omega(\frac{n}{k+1})$ if H is a spanning subgraph of M of characteristic k whose skeleton is a coarser mesh and whose branches are either of the trees in Figures 2 and 3.[7] Indeed, $\kappa_f(M,H) = \Theta(\frac{n}{k+1})$ by Lemma 3.2, since there is an embedding $\varphi : M \hookrightarrow H$ for which $d_\varphi(M,H) \leq O((\frac{n}{k+1})^{1/2})$ and $c_\varphi(M,H) \leq O((\frac{n}{k+1})^{1/2})$.

The following lemma extends the ideas above to provide a general technique for proving a lower bound on the condition number.

LEMMA 5.1 (TOPOLOGICAL LEMMA). *Let $G = (V, E_G)$ be a bounded-degree graph and $H = (V, E_H)$ be a spanning subgraph of G. Suppose there exist sets $A, B \subseteq V$ and vines $C_1, \ldots, C_r$ in H (with stems $S_1, \ldots, S_r$ and ends $\{u_1, v_1\}, \ldots, \{u_r, v_r\}$) such that: (i) $V = A \cup B \cup C_1 \cdots \cup C_r$; (ii) $A \cap B = \emptyset$, $C_i \cap C_j = \emptyset$, $A \cap C_i = \{u_i\}$, and $B \cap C_i = \{v_i\}$; and, (iii) S_i contains at least $s+1$ nodes including u_i and v_i. If the edges from a A to B contains p edges in G and q edges in H (not counting the edges in the C_i), then $\kappa_f(G,H) \geq \frac{ps}{qs+r}$.*

PROOF. The statement of the lemma is more subtle than the proof. We construct a potential vector x on the nodes for which G burns at least p units and H burns at most $q + r/s$ units. First, let the potential at the nodes of A be set to 0 and the potential at the nodes of B to 1. Next, set the potentials at the nodes along S_i, starting at u_i and ending at v_i, to $0, 1/s, \ldots, s/s$. Finally, float the potential at each remaining node in C_i to the potential of the closest node in S_i. In G, x burns at least $p \cdot 1^2$ across the cut between A and B. In H, x burns at most $q \cdot 1^2$ across the cut between A and B and $rs \cdot 1/s^2$ collectively on the C_i, for a total of at most $q + r/s$. It follows that $\kappa_f(G,H) \geq \frac{ps}{qs+r}$. □

Lemma 5.1 enables us to bound the condition number from below for the fractals we defined earlier. The proof is a simple extension of Theorem 4.1.

THEOREM 5.2. *There exists a bounded-degree planar fractal $F = (V, E_F)$ on n nodes of which every spanning subgraph $H = (V, E_H)$ of characteristic k yields $\kappa_f(F,H) \geq \Omega((\frac{n}{k+1})^{1-o(1)})$.*

PROOF. Let F be defined as in Theorem 4.1. Define the subgraph S, the cycle C, and the edge (x, y) as before. Since $H|S$ is a tree, there is a loose vine W in H whose stem contains at least $1/4$ of the nodes of C and whose loose and fixed ends, respectively, are one of either x or y (say, x) and some other node z in C. Let $A = \{x\}$, $B = V \setminus W \cup \{z\}$, and $C_1 = W$. Then we can apply Lemma 5.1 with $p = 1$, $q = 0$, $r = 1$, and $s \geq \Omega((\frac{n}{k+1})^{1-o(1)})$ to get $\kappa_f(F_i,H) \geq \Omega((\frac{n}{k+1})^{1-o(1)})$. □

Lemma 5.1 also enables us to bound the condition number from below for an expander. The proof is a simple extension of Theorem 4.2.

THEOREM 5.3. *Let $X = (V, E_X)$ be a bounded-degree expander on n nodes and $H = (V, E_H)$ be a spanning subgraph of X of characteristic k. Then $\kappa_f(X,H) \geq \Omega(\frac{n}{k+1})$.*

[7]These H are precisely the combinatorial preconditioners generated by Joshi [12] and Spielman & Teng [19] for square meshes.

PROOF. Define the edge $e = (x, y)$ and the set U as in Theorem 4.2. Let $A = U$, $B = V \setminus U$, and $C_1 = \{x, y\}$. Then we can apply Lemma 5.1 with $p \geq \Omega(n)$, $q \leq k+1$, $r = 1$, and $s = 1$ to get $\kappa_f(X,H) \geq \Omega(\frac{n}{k+1})$. □

5.2 Square Meshes

We now show our main result: a nearly-linear lower bound on the condition number for a square mesh, from which it follows that a square mesh is as hard as any graph to precondition.

Before we begin we will need a few definitions. Let $G = (V, E)$ be a graph. An ***i*-sphere** centered at vertex x are all the vertices in G at distance i from x. An ***i*-shell** centered at vertex x are all the edges in G with one endspoint in i-sphere and the other in either i-sphere or $i+1$-sphere. We state a few facts about sphere and shells true for the mesh. These seem to be the only hypothesis needed for the lower bound to follow.

LEMMA 5.4. *Let $M = (V, E_M)$ be a square mesh on n nodes. If $i \leq \sqrt{n}/2$ and Sph_i an i-sphere and Sh_i is an i-shell then*

- $|Sph_i| \geq i$
- *If a subgraph H cotains a point of Sph_i then either H contains Sph_i and Sph_{i+1} or the boundary of H contains an edge in Sh_i*
- *If $U \subseteq V$ the diameter of U is at least $\sqrt{|U|}$.*

LEMMA 5.5. *Let $M = (V, E_M)$ be a square mesh on n nodes and $H = (V, E_H)$ be a spanning subgraph of M. Suppose there is a vine $U \subset V$ in H such that the distance between the endpoints P is $2d$ in M and number of vertices of U is at most $d^2/2$ then there exists a set of boundary edges B such that*

$$\sum_{e \in B} \delta(e, P)^2 > d^3/24$$

where $\delta(x, P)$ the the distance in M from x to P in M

PROOF. Suppose the x and y are the ends of the vine U then the sheres and shell about x and y are disjoint for $1 \leq i < d$. We shall count the number of shells contains a boundary edge. We consider the shells of x and argue those of y be symetry. By Lemma 5.4 we know that any i-shell not containing a boundary edge must have its i-sphere contained in U. This gives us the inequality:

$$\sum_{i \in C} i \leq d^2/4$$

where C and $\bar{C}$ are the indices of full and not full spheres is U. Subject to 5.2 we would like to minimize:

$$\sum_{i \in \bar{C}} i^2$$

We minimize 5.2 subject to 5.2. It follows that 5.2 is minimized when $C = \{k, \ldots, d\}$ for some integer k. The lefthand side of 5.2 is at most $(d^2 - k^2)/2$ and thus $d^2/2 \leq k^2$. Thus 5.2 is mimimized when $\bar{C} = \{1, \ldots, d^2/2\}$. This gives our bound. □

We need a slightly mesier version of Lemma 5.5. The proof is essentially the same.

LEMMA 5.6. *Let $M = (V, E_M)$ be a graph statisfing the concultion of Lemma 5.4 on n nodes and $H = (V, E_H)$ be a spanning subgraph of M. Suppose there is a vine U with ends x and y. the stem from x to y contains to points x' and y' in that order such that x' and y' decompose U into three subvines U_L with ends x and x', U_M with ends x' and y', and U_R with ends y' and y with the follow properties:*

- $diaU_L$, $diaU_L \leq d/t$
- $|U_M| \leq d^2/c$
- $\delta(x, x'), \delta(y', y) \geq s$
- $\delta(x', y') \geq d$

then there exists a set of edges B from U_M to $\bar{U}$ such that

$$\sum_{e \in B} \delta(e, P)^2 \geq d \cdot s^2 / con$$

For some constant con.

PROOF. As in the last proof we consider the spheres and shells of x and y of radius $\{s+1, \ldots, s+d/2\}$. Again we only conside those centered at x. Let C be the index of the full spheres in □

LEMMA 5.7 (ISOPERIMETRIC LEMMA). *Let $M = (V, E_M)$ be a square mesh on n nodes and $H = (V, E_H)$ be a spanning subgraph of M. Suppose there is a vine $U \subset V$ in H of diameter at least d in M. Then $\kappa_f(M, H) \geq \Omega(\frac{d^3}{|U|})$.*

PROOF. We may assume that the stem of U has diameter $\frac{d}{3}$ in M; otherwise, there is another vine in H with this property. We construct a potential vector x as follows. Let the potential at every node outside U be set to ground, and set the potential at each node inside U to its distance in $H|U$ from $\partial_H U$. In H, x burns at most 1^2 across each edge in $H|U$, for a total of at most $4|U|$. In M, x burns at least $(\frac{d}{9})^2$ across each of the at least $\frac{d}{9}$ edges in $\partial_M U$ that are at distance at least $\frac{d}{9}$ from $\partial_H U$, for a total of at least $(\frac{d}{9})^3$. Hence, $\kappa_f(M, H) \geq \Omega(\frac{(d/9)^3}{4|U|})$. □

LEMMA 5.8 (TRADEOFF LEMMA). *Let $M = (V, E_M)$ be a square mesh on n nodes and $H = (V, E_H)$ be a spanning subgraph of M. Suppose there is a vine $U \subset V$ in H of diameter at least d in M. Then for any positive integer $s \leq d/3$, either: (i) $\kappa_f(M, H) \geq \Omega(ds)$; or, (ii) there is a vine $U' \subset U$ in H of diameter at least $d/15$ in M such that $|U'| \leq \frac{15s}{d}|U|$.*

PROOF. We may assume that the distance between the stem endpoints of U has $\frac{d}{3}$ in M; otherwise, there is a subvine of U with this property. Consider the finite sequence of nodes traversed by while walking from one end of the stem of U to the other. From this sequence, we can pick out a subsequence $\{w_k\}_{k=0}^{d/3s}$ of nodes, any pair of which is separated by distance at least s in M. Intuitively, this partitions the nodes in the stem of U into intervals $[w_{k-1}, w_k)$ of length at least s in M. For each k between 1 and $d/3s$, let W_k be the set of nodes in U whose closest node in the stem of U lies in the interval $[w_{k-1}, w_k)$. The (disjoint) union of the W_k partitions U.

Suppose that at least one of the $W_1, \ldots, W_{d/15s}$ and at least one of the $W_{4d/15s+1}, \ldots, W_{d/3s}$ have diameter at most $d/15s$ in M; let these be W_i and W_j, respectively. Define $C_1 = W_i \cup \{w_i\}$, $C_2 = W_j \cup \{w_{j-1}\}$, $A = W_{i+1} \cup \cdots \cup W_{j-1}$, and $B = V \setminus (A \cup C_1 \cup C_2) \cup \{w_{i-1}, w_j\}$. Then we can apply Lemma 5.1 with $p \geq d/15$, $q = 0$, $r = 2$, and s as given to get $\kappa_f(M, H) \geq \frac{1}{30} ds$.

Suppose instead that either all of the $W_1, \ldots, W_{d/15s}$ or all of $W_{4d/15s}, \ldots, W_{d/3s}$ have diameter at least $d/15s$ in M. Then trivially, one of them must enclose a vine $U' \subset U$ in H of diameter at least $d/15$ in M such that $|U'| \leq \frac{15s}{d}|U|$. □

Lemmas 4.3, 5.7, and 5.8 enable us to prove the condition number lower bound for square meshes.

THEOREM 5.9. *Let $M = (V, E_M)$ be a square mesh on n nodes and $H = (V, E_H)$ be a spanning subgraph of M of characteristic k. Then $\kappa_f(M, H) \geq \Omega((\frac{n}{k+1})^{1-o(1)})$.*

PROOF. Fix the constant $\epsilon > 0$ arbitrarily. Define $i = \lceil 1/2\epsilon \rceil$, $d = (\frac{n}{k+1})^{1/2}$ and $s = d^{1-1/i}$. By Theorem 4.3, there exists a vine $U_1 \subseteq V$ in H of size $\Omega(\frac{n}{k+1})$. The diameter of U_1 in M is $\Omega(d)$, by the isoperimetric inequality in two dimensions. By Lemma 5.8, either $\kappa_f(M, H) \geq \Omega(ds) = \Omega((\frac{n}{k+1})^{1-\epsilon})$, or there is a vine $U_2 \subset U_1$ in H of diameter $\Omega(d/15)$ in M satisfying $|U_2| \leq 15 d^{-1/i}|U_1|$. (We may assume the latter.) Applying Lemma 5.8 again, either $\kappa_f(M, H) \geq \Omega((d/15)s) = \Omega((\frac{n}{k+1})^{1-\epsilon})$, or there is a vine $U_3 \subset U_2$ in H of diameter $\Omega(d/15^2)$ in M satisfying $|U_3| \leq 15^2 d^{-2/i}|U_1|$. (Again, we may assume the latter.) We continue this process until we obtain a vine $U_i \subset U_{i-1} \subset \cdots \subset U_1$ in H of diameter $\Omega(d/15^{i-1}) = \Omega(d)$ in M and size $|U_i| \leq 15^{i-1} d^{-(1-1/i)}|U_1| \leq O((\frac{n}{k+1})^{1/2+\epsilon})$. Then by Lemma 5.7, we must have $\kappa_f(M, H) \geq \Omega((\frac{n}{k+1})^{1-\epsilon})$. Since this result holds for any sufficiently small ϵ, we conclude that $\kappa_f(M, H) \geq \Omega((\frac{n}{k+1})^{1-o(1)})$ □

6. CONCLUSION

We have demonstrated that a square mesh is as hard as any graph to precondition by a spanning tree or subgraph of prescribed Euler characteristic. It remains to construct a method to identify, based on geometric and topological properties, precisely how well a particular graph can be preconditioned and what the preconditioner should be.

In particular, we conjecture the following: if G is a hypercubic mesh in $d \geq 2$ dimensions, then there is a spanning subgraph H of G which is an (asymptotically) optimal spanner, filler, and combinatorial preconditioner of characteristic k *simultaneously*, and it satisfies $\kappa_f(G, H) = \Theta(\min_\varphi c_\varphi(G, H) \cdot d_\varphi(G, H)) = \Theta((\frac{n}{k+1})^{1-1/d} \cdot (\frac{n}{k+1})^{1/d}) = \Theta(\frac{n}{k+1})$.

7. ACKNOWLEDGMENTS

Thanks to Anupam Gupta and Bruce Maggs for pointing us to literature on dilation and congestion, respectively.

8. REFERENCES

[1] N. Alon, R. Karp, D. Peleg, and D. West. A Graph-Theoretic Game and its Application to the k-Server Problem. *SIAM Journal on Computing*, 24(1):78-100, 1995.

[2] M. Bern, J. Gilbert, B. Hendrickson, N. Nguyen, and S. Toledo. Support-Graph Preconditioners. Submitted to *SIAM Journal on Matrix Analysis and Applications*.

[3] M. Bienkowski, M. Korzeniowski, and H. Räcke. A Practical Algorithm for Constructing Oblivious Routing Schemes. In *Proceedings of the 15th ACM Symposium on Parallelism in Algorithms and Architectures*, 24-33, 2003.

[4] E. Boman, D. Chen, B. Hendrickson, and S. Toledo. Maximum-Weight-Basis Preconditioners. To appear in *Numerical Linear Algebra with Applications*.

[5] E. Boman and B. Hendrickson. Support Theory for Preconditioning. Submitted to *SIAM Journal on Matrix Analysis and Applications*.

[6] E. Boman and B. Hendrickson. On Spanning Tree Preconditioners. Manuscript, Sandia National Laboratory, 2001.

[7] K. Gremban. Combinatorial Preconditioners for Sparse, Symmetric, Diagonally Dominant Linear Systems. Ph.D. Thesis, Carnegie Mellon University, 1996.

[8] K. Gremban, G. Miller, and M. Zagha. Performance Evaluation of a New Parallel Preconditioner. In *Proceedings of the 9th IEEE International Parallel Processing Symposium*, 65-69, 1995.

[9] A. Gupta. Steiner Points in Tree Metrics Don't (Really) Help. In *Proceedings of the 12th ACM/SIAM Symposium on Discrete Algorithms*, 220-227, 2001.

[10] A. Gupta, R. Krauthgamer, and J. Lee. Bounded Geometries, Fractals, and Low-Distortion Embeddings. In *Proceedings of the 44th IEEE Symposium on Foundations of Computer Science*, 534-543, 2003.

[11] C. Harrelson, K. Hildrum, and S. Rao. A Polynomial-Time Tree Decomposition to Minimize Congestion. In *Proceedings of the 15th ACM Symposium on Parellelism in Algorithms and Architectures*, 34-43, 2003.

[12] A. Joshi. Topics in Optimization and Sparse Linear Systems. Ph.D. Thesis, University of Illinois at Urbana-Champaign, 1997.

[13] T. Leighton, B. Maggs, and S. Rao. Packet Routing and Job-Shop Scheduling in $O(\text{Congestion} + \text{Dilation})$ Steps. *Combinatorica*, 14(2):167-186, 1994.

[14] B. Maggs, G. Miller, O. Parekh, R. Ravi, and S.-L. Woo. M-Matrices Have Good Support Tree Preconditioners. Manuscript, 2002.

[15] K. Naidu and H. Ramesh. Lower Bounds for Embedding Graphs into Graphs of Smaller Characteristic. In *Proceedings of the 22nd Conference on Foundations of Software Technology and Theoretical Computer Science*, 301-310, 2002.

[16] Y. Rabinovich and R. Raz. Lower Bounds on the Distortion of Embedding Finite Metric Spaces in Graphs. *Discrete and Computational Geometry*, 19:79-94, 1998.

[17] H. Räcke. Minimizing Congestion in General Networks. In *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science*, 43-52, 2002.

[18] J. Reif. Efficient Approximate Solution of Sparse Linear Systems. *Computers and Mathematics with Applications*, 36(9):37-58, 1998.

[19] D. Spielman and S.-H. Teng. Solving Sparse, Symmetric, Diagonally Dominant Linear Systems in Time $O(m^{1.31})$. In *Proceedings of the 44th IEEE Symposium on Foundations of Computer Science*, 416-427, 2003.

[20] D. Spielman and S.-H. Teng. Nearly-Linear Time Algorithms for Graph Partitioning, Graph Sparsification, and Solving Linear Systems. To appear in *Proceedings of the 36th ACM Symposium on Theory of Computing*, 2004.

[21] P. Vaidya. Solving Linear Equations with Symmetric Diagonally Dominant Matrices by Constructing Good Preconditioners. Manuscript, 1990.