Journal of Combinatorial Optimization, 1:79–104 (1997)

Moments of Inertia and Graph Separators

KEITH D. GREMBAN*
Advanced Information Systems Division, CTA Inc. 5670 Greenwood Plaza Blud. Suite 200, Englewood, CO 80111. kgremban@ctaeng.com

GARY L. MILLER*
School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, glmiller@cs.cmu.edu

SHANG-HUA TENG†
Department of Computer Science, University of Minnesota, Minneapolis, MN 55455, steng@cs.umn.edu

(Received 6 May 1996; Accepted 15 May 1996)

Abstract

Graphs that arise from the finite element or finite difference methods often include geometric information such as the coordinates of the nodes of the graph. The geometric separator algorithm of Miller, Teng, Thurston, and Vavasis uses some of the available geometric information to find small node separators of graphs. The algorithm utilizes a random sampling technique based on the uniform distribution to find a good separator. We show that sampling from an elliptic distribution based on the inertia matrix of the graph can significantly improve the quality of the separator. More generally, given a cost function f on the unit d-sphere U_d, we can define an elliptic distribution based on the second moments of f. The expectation of f with respect to the elliptic distribution is less than or equal to the expectation with respect to the uniform distribution, with equality only in degenerate cases. We also demonstrate experimentally that the benefit gained by the use of the additional geometric information is significant. Some previous algorithms have used the moments of inertia heuristically, and suffer from extremely poor worst case performance. This is the first result, to our knowledge, that incorporates the moments of inertia into a provably good strategy.

1. Introduction

Many problems in computational science and engineering are based on unstructured meshes of points in two or three dimensions. The meshes can be quite large, often containing millions of points. Typically, the size of the mesh is limited by the size of the machine available to solve the problem, even though, in many problems, the accuracy of the solution is related to the size of the mesh. As a result, methods for solving problems on large meshes are becoming increasingly important.

*Supported in part by NSF Grant CCR-9016641 and CCR-9505472.
†Supported in part by an NSF CAREER award (CCR-9502540) and a Sloan Fellowship. Part of the work was done while the author was at Department of Mathematics and Lab. for Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

One method is to partition the mesh and assign the pieces to separate processors of a parallel machine. For maximum efficiency, each processor should be responsible for pieces of roughly the same size, and communication between pieces should be minimized. Mesh partitioning is the process of decomposing a mesh into two or more pieces of roughly equal size. A mesh consists of nodes (vertices) and undirected edges connecting the nodes. In some cases, additional information in the form of the geometric coordinates of the vertices may also be available. Meshes are special cases of graphs, and mesh partitioning is a special case of the more general problem of graph partitioning.

The graph partitioning problem can be defined as follows. Given a graph $G = (V, E)$ where V is the set of vertices of G and E is the set of edges of G, find a set partition $A \cup B = V$ such that the size of A is approximately equal to that of B and the number of edges between A and B is small. The set of edges between A and B is known as an edge separator. A collection of vertices whose removal induces a partition of the graph is known as a vertex separator. The goal of graph partitioning is to find small separators.

In this paper, we shall be concerned with vertex separators, which are formalized in the following definition.

Definition 1.1. (Vertex Separators) A subset of vertices C of a graph G with n vertices is an $f(n)$-separator that δ-splits if $|C| \leq f(n)$ and the vertices of G–C can be partitioned into two sets A and B such that there are no edges from A to B, $|A|, |B| \leq \delta n$, where f is a function and $0 \leq \delta \leq 1$.

Not every graph has a small separator; for example, consider the complete graph on n points. A significant amount of research has gone towards answering the question of which families of graphs have small separators.

Two of the most well-known families of graphs that have small separators are trees and planar graphs. Every tree has a single vertex separator that 2/3-splits [22]. Lipton and Tarjan [28] proved that every planar graph has a $\sqrt{8n}$-separator that 2/3-splits. Their result improved an earlier one by Ungar [37]. Some extensions of their work have been made [5, 12, 13, 29]. Gilbert, Hutchinson, and Tarjan showed that all graphs with genus bounded by g have an $O(\sqrt{gn})$-separator [16], and Alon, Seymour, and Thomas proved that all graphs with an excluded minor isomorphic to the h-clique have an $O(h^{3/2}\sqrt{n})$-separator [1].

Graph partitioning is the basis for a number of techniques for solving problems involving large graphs. For example: solution of a finite element problem on a distributed memory parallel processor requires partitioning the graph to assign roughly equal numbers of nodes to each processor, while minimizing the communication requirements between processors [2, 9, 32, 35, 36, 38]; efficient node ordering for solving linear systems is related to finding good partitions [15, 27, 33]; optimizing the physical layout of the components of a VLSI circuit involves graph partitioning [6, 24, 26].

Many different approaches to graph partitioning have been developed. Graph partitioning algorithms can be classified as being either: *combinatorial* or *geometric*. A combinatorial partitioning algorithm only makes use of the graph connectivity. Examples of combinatorial partitioning algorithms include: iterative improvement [23, 20], simulated

annealing [32, 38], spectral partitioning [7, 19, 21, 34, 35], the greedy method [9], and multicommodity flow [25]. Geometric approaches to graph partitioning make use of the spatial coordinates of the vertices of the graph. Examples of geometric partitioning algorithms include: recursive coordinate bisection [18, 35, 38], inertia-based slicing algorithms [9], and the sphere separator algorithm [30, 31]. However, all of the above approaches, excluding the multicommodity flow algorithm of Leighton and Rao [25] and the sphere separator algorithm of Miller, Teng, Thurston and Vavasis [30], fail to provide any performance guarantees.

In this paper, we explore an extension to the sphere separator algorithm of Miller, Teng, Thurston, and Vavasis [30]. In particular, we show that the moments of inertia of the population of graph vertices can be used to improve a random sampling technique employed in the algorithm. Our improved method has a provably good bound on well-shaped meshes and geometric graphs.

Some previous graph partitioning algorithms have used the moments of inertia heuristically, and suffer from extremely poor worst case performance. In contrast, we present a technique that takes advantage of the moments of inertia, while avoiding the problems that can lead to poor performance. We believe that this is the first result to show how moments of inertia can be incorporated into a provably good strategy. While we apply moments of inertia to the problem of graph partitioning, the technique is generalizable to other problems.

This paper is organized as follows. First, we present an overview of geometric graph partitioning and a brief, intuitive explanation of the original sphere separator algorithm. Next, we present an abstract version of graph partitioning and state two related optimization problems. We then prove that the expected cost of sampling from an elliptic distribution based on the moments of inertia is less than that from a uniform distribution; this is first proved in one dimension, and then extended to higher dimensions. Then, we show how this result can be incorporated into the sphere separator algorithm. Finally, we present experimental results that demonstrate that dramatic improvement that can result from making use of the moments of inertia to partition graphs.

2. Geometric partitioning

Geometric partitioning is only applicable to the class of graphs for which geometric information such as the coordinates of the vertices is known. Most of the geometric partitioning algorithms are heuristic in the sense that no bounds on performance have been proved. The sphere separator algorithm is distinct in that bounds on both the expected running time and the expected quality of the separator can be proved.

2.1. Recursive coordinate bisection

The simplest form of geometric partitioning is recursive coordinate bisection (RCB) [35, 38]. In the RCB algorithm, the vertices of the graph are projected onto one of the

coordinate axes, and the vertex set is partitioned around a hyperplane through the median of the projected coordinates. Each resulting subgraph is then partitioned along a different coordinate axis until the desired number of subgraphs is obtained.

Because of the simplicity of the algorithm, RCB is very quick and cheap, but the quality of the resultant separators can vary dramatically, depending on the embedding of the graph in $\mathbb{R}^d$. For example, consider a graph that is "+"-shaped. Clearly, the best (smallest) separator consists of the vertices lying along a diagonal cut through the center of the graph. RCB, however, will find the largest possible separators, in this case, planar cuts through the centers of the horizontal and vertical components of the graph.

2.2. Inertia-based slicing

Williams [38] noted that RCB had poor worst case performance, and suggested that it could be improved by slicing orthogonal to the principal axes of inertia, rather than orthogonal to the coordinate axes. Farhat and Lesoinne implemented and evaluated this heuristic for partitioning [9].

In three dimensions, let $v = (v_x, v_y, v_z)^t$ be the coordinates of vertex v in $\mathbb{R}^3$. Then the inertia matrix I of the vertices of a graph with respect to the origin is given by

$$I = \begin{pmatrix} I_{xx} & I_{xy} & I_{xz} \\ I_{yx} & I_{yy} & I_{yz} \\ I_{zx} & I_{zy} & I_{zz} \end{pmatrix}$$

where,

$$I_{xx} = \sum_{v \in V} v_y^2 + v_z^2, \quad I_{yy} = \sum_{v \in V} v_x^2 + v_z^2, \quad I_{zz} = \sum_{v \in V} v_x^2 + v_y^2$$

and, for $i, j \in \{x, y, z\}$, $i \neq j$,

$$I_{ij} = I_{ji} = -\sum_{v \in V} v_i v_j$$

The eigenvectors of the inertia matrix are the principal axes of the vertex distribution. The eigenvalues of the inertia matrix are the principal moments of inertia. Together, the principal axes and principal moments of inertia define the inertia ellipse; the axes of the ellipse are the principal axes of inertia, and the axis lengths are the square roots of the corresponding principal moments. Physically, the size of the principal moments reflect how the mass of the system is distributed with respect to the corresponding axis—the larger the principal moment, the more mass is concentrated at a distance from the axis.

Let I_1, I_2, and I_3 denote the principal axes of inertia corresponding to the principal moments $\alpha_1 \leq \alpha_2 \leq \alpha_3$. Farhat and Lesoinne projected the vertex coordinates onto I_1, the axis about which the mass of the system is most tightly clustered, and partitioned using a

planar cut through the median. This method typically yielded a good initial separator, but did not perform as well recursively on their test mesh—a regularly structured "T"-shape.

Farhat and Lesoinne did not present any results on the theoretical properties of the inertia slicing method. In fact, there are pathological cases in which the inertia method can be shown to yield a very poor separator. Consider, for example, a "+"-shape in which the horizontal bar is very wide and sparse, while the vertical bar is relatively narrow but dense. I_1 will be parallel to the horizontal axis, but a cut perpendicular to this axis through the median will yield a very large separator. A diagonal cut will yield the smallest separator, but will not be generated with this method.

2.3. The sphere separator algorithm

The geometric partitioning algorithms outlined above both employ heuristics for finding separators, and utilize cuts defined by hyperplanes through the graph. Pathological cases were described for which the heuristics would yield extremely poor performance. The sphere separator algorithm is unique in that it is not heuristic; proofs exist for bounds on expected performance. Additionally, it can be shown that planar cuts are not sufficiently robust to yield good separators for all types of graphs; the sphere separator algorithm uses spherical cuts, which can be shown not to suffer from the inadequacies of planar cuts [36].

This subsection presents an intuitive overview of the sphere separator algorithm. The full algorithm can be found in [17, 30]; the relevant theoretical background is contained in [3, 8, 17, 36]. The following definition is needed.

Definition 2.1. (Center Points) Let P be a finite set of points in $\mathbb{R}^d$. For each $0 < \delta < 1$, a point $c \in \mathbb{R}^d$ is a δ-center point of P if every hyperplane containing c δ-splits P.

The basic algorithm is stated in Figure 1. In the basic algorithm, stereo-up is the

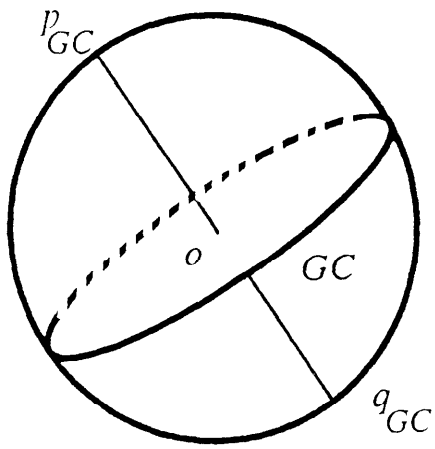

Figure 1. The basic sphere separator algorithm.

standard stereographic projection mapping which can be described as follows. Assume the graph is embedded in $\mathbb{R}^d$ coordinate plane, and let U_d be the sphere in $\mathbb{R}^{d+1}$ centered at the origin. For each vertex v in the graph, construct the line passing through v and the north pole of U_d. The line intersects U_d at q, which is the stereographic projection of v.

Essentially, the algorithm operates by finding a conformal mapping of the set of vertices V of the graph in $\mathbb{R}^d$ onto U_d such that origin is a center point, and the "mass" of V is spread out more-or-less evenly over the surface of U_d, in the sense that every hyperplane through the origin of U_d partitions V into two sets of roughly equal size. Since any hyperplane through the origin defines a great circle on U_d, an equivalent statement is that any great circle of U_d partitions V appropriately. Any great circle therefore defines a sphere separator S in $\mathbb{R}^d$. The vertices of the graph that lie "near" S then comprise a vertex separator of G. Because any great circle of U_d defines a balanced separator, the algorithm chooses great circles randomly. In Section 7, we will review the theoretical results of [30].

The only information used in the determination of the conformal mapping is the location of an approximate center point of the set of points. Hence, the mapping only approximately distributes the points evenly across the sphere. This observation is the key to using moments of inertia. As will be shown in the next section, the moments of inertia of the point set reveals important information about the distribution of the points on the sphere, and can be used to guide the selection of a great circle to reduce the size of the resultant vertex separator.

3. The abstract problem

To prove the utility of the moments of inertia, it is useful to abstract away from the discrete set of points that define a graph G, and instead consider a continuous analog. Let g: $\mathbb{R}^d \to \mathbb{R}^+$ be a density function with compact support. Let γ represent the support of g. In two dimensions, γ can be visualized as representing a sheet of material with varying density g. To make the connection to the graph G, consider that g reflects the density of the vertices in G.

The abstract partitioning problem consists of finding a cut S through γ such that γ is partitioned into two sets of approximately equal size, and the size of the cut is small. The sphere separator algorithm outlined above, transforms the abstract partitioning problem into a simpler one. The conformal mapping of γ onto U_d in effect spreads γ out, and ensures that for any great circle GC, the integrals of g over the hemispheres defined by GC are approximately equal. Thus, the problem reduces to that of finding a great circle GC such that the integral of g over GC is small.

Two related problems can now be defined.

Problem 3.1. (primal problem) Given a density function f: $U_d \to \mathbb{R}^+$, find a point u of U_d of small cost where the cost of u is $f(u)$.

Problem 3.2. (dual problem) Given a density function f: $U_d \to \mathbb{R}^+$, find a great circle C of U_d of small cost, where the cost of C is $\int_{v \in C} f(v) ds^{d-1}$, the integral along the great circle C.

4. One dimension

To illustrate the method and introduce concepts and notations, we start with one dimension. That is, assume that we are given a nonnegative real function f: $U_1 \to \mathbb{R}^+$, which is defined on the unit circle. In other words, $f(\theta)$ is a nonnegative real for an angle θ. The goal of the primal problem is to find a solution angle θ, such that $f(\theta)$ is minimal.

This problem can naturally be viewed as a mathematical optimization problem. But, in some applications, either f is not explicitly given, or the gradient of f is too expensive to compute. Further, in some cases, such as geometric separators, the average value of f is reasonably small, and the determination of the exact minimum is unnecessary.

One simple method for finding a solution θ is to choose an angle randomly from a uniform distribution. This is essentially the approach taken in the sphere separator algorithm. The expected cost (normalized by multiplying a factor of 2π) of the solution, denoted by $UC(f)$, is clearly $UC(f) = \int_0^{2\pi} f(\theta) d\theta$.

Given a random sampling approach, the question naturally arises as to whether information exists about the distribution of f that can improve the expectation. In particular, is there an efficiently computable probability distribution with an expectation less than or equal to that of the uniform distribution?

4.1. Distribution as an angle updating function

One way to describe some probabilistic distributions in one dimension is to use an angle updating function. Suppose $\delta(\theta)$ is an angle function. A new distribution defined by $\delta(\theta)$ can be constructed as follows: choose a uniform random angle θ and return $\theta + \delta(\theta)$. The expected cost (again, normalized by a factor of 2π) of the new distribution, denoted as $KC(f, \delta)$, is then given by $KC(f, \delta) = \int_0^{2\pi} f(\theta + \delta(\theta)) d\theta$.

Lemma 4.1. If $\delta(0) = \delta(2\pi) = 0$ and for all $0 \le \theta \le 2\pi$, $\theta + \delta(\theta) \le 2\pi$ then

$$\Delta_f(\delta) = UC(f) - KC(f, \delta) = \int_0^{2\pi} [f(\theta + \delta(\theta))\delta'(\theta)] d\theta.$$

Proof: We can prove the lemma by Taylor's expansion and integration by parts. The following proof is a more intuitive one. Because $\delta(0) = \delta(2\pi) = 0$ and for all $0 \le \theta \le 2\pi$, $\theta + \delta(\theta) \le 2\pi$, by changing of variables, we have

$$\begin{aligned} UC(f) = \int_0^{2\pi} f(\theta)d\theta &= \int_0^{2\pi} f(\theta + \delta(\theta))d(\theta + \delta(\theta)) \\ &= \int_0^{2\pi} f(\theta + \delta(\theta))d(\theta + \delta(\theta)) \\ &= \int_0^{2\pi} f(\theta + \delta(\theta))[d\theta + \delta'(\theta)d\theta] \\ &= [\int_0^{2\pi} f(\theta + \delta(\theta))d\theta] + [\int_0^{2\pi} f(\theta + \delta(\theta))\delta'(\theta)d\theta] \\ &= KC(f, \delta) + [\int_0^{2\pi} f(\theta + \delta(\theta))\delta'(\theta)d\theta]. \end{aligned}$$

□

4.2. Inertia based distribution by angle updating

Again, suppose $f(\theta)$ is the cost function on the unit circle. Define

$$I(f) = \int_0^{2\pi} f(\theta)d\theta$$

$$I_{xx}(f) = \int_0^{2\pi} f(\theta) \sin^2 \theta d\theta \quad I_{yy}(f) = \int_0^{2\pi} f(\theta) \cos^2 \theta d\theta$$

The quantities I_{xx} and I_{yy} are called the moments of inertia of f with respect to the x and y axes, respectively. Clearly, for all f, $I(f) = I_{xx}(f) + I_{yy}(f)$, because $\sin^2 \theta + \cos^2 \theta = 1$, and the expected cost of the uniform distribution, $UC(f)$, is equal to $I(f) = I_{xx}(f) + I_{yy}(f)$.

Let $\lambda \geq 1$ be a constant (independent of θ) to be specified later. We construct a new distribution based on λ. In particular, we define an angle updating function $\delta_\lambda(\theta)$ as following:[1] $\delta(\theta) = \tan^{-1} [\lambda \tan \theta] - \theta$. Geometrically, $\delta(\theta)$ is the angular difference of the vector $(\cos \theta, \lambda \sin \theta)$ and the vector $(\cos \theta, \sin \theta)$.

It is easy to check that $\delta(k\pi/2) = 0$ for all $k \in \{0, 1, 2, 3, \ldots\}$. So it satisfies the condition of Lemma 4.1. The equation $\cos^2 \theta + \lambda^2 \sin^2 \theta = 1$ defines an ellipse with axes of lengths 1 and $1 / \lambda$, so sampling points uniformly from the unit circle and applying the angle updating formula is equivalent to sampling points uniformly from the boundary of the corresponding ellipse.

Theorem 4.2. If $\lambda = \sqrt{I_{yy}/I_{xx}}$, then the expected cost of the distribution defined by δ_λ is equal to $2\sqrt{I_{xx}I_{yy}}$. Hence, $\Delta_f(\delta_\lambda) = (\sqrt{I_{yy}} - \sqrt{I_{xx}})^2$, which is positive as long as $I_{yy} \neq I_{xx}$.

Proof: Let $\beta = \theta + \delta(\theta)$. We have $\tan \beta = \tan(\theta + \delta(\theta)) = \lambda \tan \theta$. Hence, we have

[1]For simplicity, we use $\delta(\theta)$ instead of $\delta_\lambda(\theta)$.

$$\cos^2\theta = \frac{\cos^2\beta}{\cos^2\beta + \sin^2\beta/\lambda^2} \tag{1}$$

By differentiating both side of $\tan(\theta + \delta(\theta)) = \lambda \tan\theta$ (with respect to θ, of course), we have

$$\frac{1 + \delta'(\theta)}{\cos^2(\theta + \delta(\theta))} = \lambda \frac{1}{\cos^2\theta}. \tag{2}$$

Therefore,

$$\begin{aligned} 1 + \delta'(\theta) &= \lambda \frac{\cos^2(\theta + \delta(\theta))}{\cos^2\theta} \\ &= \lambda \frac{\cos^2\beta}{\cos^2\theta} \\ &= \lambda(\cos^2\beta + \sin^2\beta/\lambda^2) \\ &= \lambda\left(1 - \frac{\lambda^2 - 1}{\lambda^2}\sin^2\beta\right) \end{aligned}$$

We now calculate the cost discrepancy of this distribution over the uniform distribution. By Lemma 4.1, we have

$$\begin{aligned} &\Delta_f(\delta) \\ &= \int_0^{2\pi} [f(\theta) - f(\theta + \delta(\theta))]d\theta \\ &= \int_0^{2\pi} f(\theta + \delta(\theta))\delta'(\theta)d\theta \end{aligned}$$

(Change of variable, $\beta = \theta + \delta(\theta)$)

$$= \int_0^{2\pi} f(\beta)\left(1 - \frac{1}{1 + \delta'(\theta)}\right)d\beta$$

(Using Equation (3))

$$= \int_0^{2\pi} f(\beta)\left(1 - \frac{1}{\lambda}\left[\frac{1}{1 - \frac{\lambda^2 - 1}{\lambda^2}\sin^2\beta}\right]\right)d\beta$$

(Using $\frac{1}{1 - x} = 1 + \sum_{i=1}^{\infty} x^i$, for $0 < x < 1$)

$$= \int_0^{2\pi} f(\beta)\left(1 - \frac{1}{\lambda}\left[1 + \sum_{i=1}^{\infty}\left(\frac{\lambda^2 - 1}{\lambda^2}\sin^2\beta\right)^i\right]\right)d\beta$$

(Using $\sin^{2i}\theta \leq \sin^2\theta$ for $i \geq 1$)

$$\geq \int_0^{2\pi} f(\beta)\left(1 - \frac{1}{\lambda}\left[1 + \sum_{i=1}^{\infty}\left(\frac{\lambda^2 - 1}{\lambda^2}\right)^i \sin^2\beta\right]\right)d\beta$$

$$= \int_0^{2\pi} f(\beta)\left(1 - \frac{1}{\lambda}[1 + (\lambda^2 - 1)\sin^2\beta]\right)d\beta$$

$$= I_{xx} + I_{yy} - \frac{1}{\lambda}(I_{xx} + I_{yy} + (\lambda^2 - 1)I_{xx})$$

$$= I_{xx} + I_{yy} - I_{yy}/\lambda - \lambda I_{xx}$$

If $\lambda = \sqrt{I_{yy}/I_{xx}}$ (such that the last quantity is maximized), then $\Delta_f(\delta_\lambda) = (\sqrt{I_{yy}} - \sqrt{I_{xx}})^2$. □

Theorem 4.2 proves that the use of an angle updating function corresponding to an ellipse yields an improvement in the expected cost. Moreover, the axes of the ellipse are proportional to the principal moments of inertia of the cost distribution. In the next section, we generalize this result to higher dimensions.

5. Two and higher dimensions

In the previous section, we defined the moments of inertia of a cost function defined on U_1. The principal axes and principal moments of the cost function define an ellipse, which in turn defines an elliptic probability distribution. The expected cost of sampling from the elliptic distribution is less than or equal to that obtained by sampling from the uniform distribution. The moments of inertia of a function defined on U_1 have an intuitive physical interpretation, but lead to a difficult proof that does not easily generalize.

In this section, we consider the primal problem, starting initially in two dimensions, and then extending the results to higher dimensions. Rather than using moments of inertia, we instead use a dual concept—the second moments with respect to the origin. On the unit sphere in $d + 1$ dimensions, U_d, the second moments are related to the moments of inertia (see Section 6), and the use of the second moments lead to cleaner proofs.

5.1 Two dimensions

To find a sphere separator for a two-dimensional graph, we stereographically project the points of the graph onto U_2, the unit sphere in $\mathbb{R}^3$. The goal of the primal problem is to find a point $u \in U_2$ that minimizes $f(u)$, the cost of u.

One simple method for finding u is to repeatedly choose a point randomly from U_2 based on a uniform distribution. The expected cost per choice with this distribution is denoted $\mu(f)$, and is given by

$$\mu(f) = \frac{1}{|U_2|} \int_{w \in U_2} f(w) dA$$

where $|U_2|$ is the surface area of U_2. A point on the unit sphere can be uniquely represented by three parameters (α, β, γ) satisfying $\alpha^2 + \beta^2 + \gamma^2 = 1$. For example, in spherical coordinates, $\alpha(\theta, \phi) = \sin\phi \cos\theta$, $\beta(\theta, \phi) = \sin\phi \sin\theta$, $\gamma(\theta, \phi) = \cos\phi$. For each point w on the unit sphere, let $(\alpha(w), \beta(w), \gamma(w))$ be its three coordinates. Define for $i \in \{\alpha, \beta, \gamma\}$,

$$\mu_{ii}(f) = \frac{1}{|U_2|} \int_{w \in U_2} f(w)(i(w))^2 dA$$

Then, $\mu_{ii}(f)$ is the second moment of f with respect to the i axis. We call the following matrix the moment matrix of f.

$$\mathcal{M}(f) = \begin{pmatrix} \mu_{\alpha\alpha} & \mu_{\alpha\beta} & \mu_{\alpha\gamma} \\ \mu_{\beta\alpha} & \mu_{\beta\beta} & \mu_{\beta\gamma} \\ \mu_{\gamma\alpha} & \mu_{\gamma\beta} & \mu_{\gamma\gamma} \end{pmatrix}$$

Notice that the *inertia matrix* of f, $I(f)$, is defined as following:

$$I(f) = \begin{pmatrix} \mu_{\beta\beta} + \mu_{\gamma\gamma} & -\mu_{\alpha\beta} & -\mu_{\alpha\gamma} \\ -\mu_{\beta\alpha} & \mu_{\alpha\alpha} + \mu_{\gamma\gamma} & -\mu_{\beta\gamma} \\ -\mu_{\gamma\alpha} & -\mu_{\gamma\beta} & \mu_{\alpha\alpha} + \mu_{\beta\beta} \end{pmatrix}$$

Since $\alpha^2 + \beta^2 + \gamma^2 = 1$ and integration is linear, we have that

$$\mu(f) = \mu_{\alpha\alpha}(f) + \mu_{\beta\beta}(f) + \mu_{\gamma\gamma}(f).$$

Lemma 5.1. For all density functions f on U_2,

$$\mathcal{M}(f) + I(f) = \mu(f) I_3,$$

where I_d denotes the $d \times d$ identity matrix.

5.2. Elliptic distribution

Uniform sampling of points on U_1, the unit circle, is equivalent to sampling angles θ uniformly from $[0, 2\pi)$. An identical distribution is obtained by picking a point p uni-

formly from B_2, the unit ball, and if $p \neq 0$ normalizing p to unit length, i.e., projecting p onto U_1. If we are so unlucky as to pick $p = 0$ then we repeat the process until we get a $p \neq 0$.

We define our elliptic distribution, denoted dist(E), with respect to an ellipse or ellipsoid E of dimension d and unit sphere U_{d-1} via the following random process:

Process 5.2. While $p = 0$ pick a random point $p \in E$ and if $p \neq 0$ return the projection of p onto U_{d-1}.

As an example consider the ellipse E_{ab} defined by

$$E_{ab} = \{(x, y) \mid x^2/a^2 + y^2/b^2 \leq 1.\}$$

Clearly, the sample is biased by the shape of the ellipse. The associated probability density function is given by the following lemma:

Lemma 5.3. The probability density function for dist(E_{ab}) over U_1 is

$$\frac{1}{2\pi ab}\left[\frac{1}{x^2/a^2 + y^2/b^2}\right]$$

Proof: Let $p = (x, y)$ be a point on U_1. Further let L be the distance from 0 through p to the boundary of E_{ab}, i.e., a scalar $L \geq 0$ such that the point $L \cdot p = (Lx, Ly)$ lies on the boundary of E_{ab}. Thus L satisfies $L^2x^2/a^2 + L^2y^2/b^2 = 1$. For each small line segment ds on U_1 we must determine the area of E_{ab} mapped to ds. In the limit as ds goes to zero this will be a triangle of area $(1/2)L^2 ds$. Now the integral over U_1 will give us the area of the ellipse, i.e.:

$$|E_{ab}| = \int_{x \in U_1} \frac{L^2(x)}{2}\, ds$$

Thus the probability density is $L^2/(2|E_{ab}|)$. Using the fact that $|E_{ab}| = \pi ab$ the Lemma follows □

We next compute the probability density function for the distribution dist(E_{abc}) where E_{abc} is the ellipsoid given by:

$$E_{abc} = \{(x, y, z) \mid x^2/a^2 + y^2/b^2 + z^2/c^2 \leq 1\}.$$

Lemma 5.4. The probability density function for dist(E_{abc}) over U_2 is, where $|B_3|$ is the volume of B_3,

$$\frac{1}{3|B_3|abc}\left(\frac{1}{x^2/a^2 + y^2/b^2 + z^2/c^2}\right)^{3/2}$$

Proof: Let ds^2 be a small square patch on U_2. As in the proof of Lemma 5.3 we will determine that part of E_{abc} which is projected onto U_2. Normalizing this volume will then give us our probability density function.

Let $p = (\alpha, \beta, \gamma)$ be a point on U_2 in the patch ds^2 and L be the distance from from 0 through p to the boundary of E_{abc}. It follows that $L = (\alpha^2/a^2 + \beta^2/b^2 + \gamma^2/c^2)^{-1/2}$. The preimage of ds^2 is a pyramid shaped object as ds goes to zero with volume $L^3 / 3\ ds^2$. As in Lemma 5.3, the integral over U_2 of the unnormalized density function is:

$$|E_{abc}| = \int_{p \in U_2} \frac{L^3(p)}{3}\, ds^2$$

Thus, the probability density is $L^3/(3|\ E_{abc}|)$. Using the fact that $|\ E_{abc}| = (4/3)\pi abc = |\ B_3|\ abc$ the Lemma follows □

Before we estimate the expectation of dist(E_{abc}) we will give some inequalities to simplify the integrals. We will use the following lemma in several of the proofs.

Lemma 5.5. Suppose a, b, α, β are nonnegative reals, and $\alpha + \beta \leq 1$. Then $(a\alpha + b\beta)^k \leq a^k\alpha + b^k\beta$.

Proof: We prove it by induction on k. The lemma is clearly true for $k = 1$. Now assume it is true for $k - 1$, we prove for k.

$$\begin{aligned}
&(a\alpha + b\beta)^k \\
&= (a\alpha + b\beta)^{k-1}\,(a\alpha + b\beta) \\
&\leq (a^{k-1}\,\alpha + b^{k-1}\,\beta)\,(a\alpha + b\beta) \\
&= a^k\,\alpha + b^k\,\beta \\
&\quad - [a^k\,\alpha(1 - \alpha) + b^k\,\beta(1 - \beta)] \\
&\quad - [a^{k-1}\,b\alpha\beta + ab^{k-1}\,\alpha\beta] \\
&\leq a^k\,\alpha + b^k\,\beta - \alpha\beta(a^{k-1} - b^{k-1})\,(a - b) \\
&\leq a^k\,\alpha + b^k\,\beta
\end{aligned}$$

□

Lemma 5.6. Let $\alpha^2 + \beta^2 + \gamma^2 = 1$ then

$$\left(\frac{1}{\alpha^2/a^2 + \beta^2/b^2 + \gamma^2/c^2}\right)^{3/2} \leq a^3\,\alpha^2 + b^3\,\beta^2 + c^3\,\gamma^2$$

Proof: Here we consider the special case when $c = 1$. Substituting $1 - \alpha^2 + \beta^2$ for γ^2 and rearranging terms we get that the LHS equals $(1/(1 - (u\alpha^2 + v\beta^2)))^{3/2}$ where $u = (a^2 - 1)/a^2$ and $v = (b^2 - 1)/b^2$. Since $u\alpha^2 + v\beta^2 < 1$ we can expand $1/(1 - (u\alpha^2 + v\beta^2)$ as Taylor series about zero. Since the constant term in this series is one we can expand the square root of this series. There are two expansions and we take the positive one, with constant term one. We now take the cube of this series. This gives a series of the form below with $t_0 = 1$. Using Lemma 5.5,

$$\sum_{i=0}^{\infty} t_i(u\alpha^2 + v\beta^2)^i$$

$$\leq \sum_{i=0}^{\infty} t_i(u^i\,\alpha^2 + v^i\,\beta^2)$$

$$\leq \sum_{i=0}^{\infty} t_i u^i\,\alpha^2 + \sum_{i=0}^{\infty} t_i v^i\,\beta^2 + (1 - \alpha^2 - \beta^2)$$

$$= (1/(1 - u))^{3/2}\,\alpha^2 + (1/(1 - v))^{3/2}\,\beta^2 + \gamma^2$$

$$= a^3\,\alpha^2 + b^3\,\beta^2 + \gamma^2$$

By the way we generated the t_is we get the second equality. □

We let *Expect*(E, f) denote the expected cost (normalized by $3|B_3|$) when sampling from the elliptic distribution given by ellipsoid E and using cost function f. Now the expectation of the distribution of dist(E_{abc}, f) is given by the following integral:

$$\text{Expect}(E_{abc}, f) = (3|B_3|)\int_{p \in U_2} f(p)\,\frac{1}{3|B_3|abc}\,L^3\,dA$$

where $L = (\alpha^2/a^2 + \beta^2/b^2 + \gamma^2/c^2)^{-1/2}$. Combining with Lemma 5.6 we get the following inequality:

Lemma 5.7. *Expect*$(E_{abc}, f) \leq (a^3\mu_{\alpha\alpha} + b^3\mu_{\beta\beta} + c^3\mu_{\gamma\gamma})/abc$

If we now minimize the right hand side of Lemma 5.7 we get the following Theorem:

Theorem 5.8. If $a = (1/\mu_{\alpha\alpha})^{1/3}$, $b = (1/\mu_{\beta\beta})^{1/3}$, and $c = (1/\mu_{\gamma\gamma})^{1/3}$, then the expected cost of sampling from the elliptic distribution is,

$$Expect(E_{abc}, f) = 3(\mu_{\alpha\alpha}\mu_{\beta\beta}\mu_{\gamma\gamma})^{1/3}$$

Moreover, by [4] $3(\mu_{\alpha\alpha}\mu_{\beta\beta}\mu_{\gamma\gamma})^{1/3} \leq \mu_{\alpha\alpha} + \mu_{\beta\beta} + \mu_{\gamma\gamma}$, we have

$$Expect(E_{abc}, f) \leq UC(f),$$

where $UC(f)$ is the normalized (by a factor $3|B_3|$) cost when the distribution is uniform.

If we divide both the uniform cost and the elliptic cost by three, then the uniform cost is the arithmetic mean of $\mu_{\alpha\alpha}$, $\mu_{\beta\beta}$, $\mu_{\gamma\gamma}$, and the elliptic cost is bounded from above by the geometric mean of $\mu_{\alpha\alpha}$, $\mu_{\beta\beta}$, $\mu_{\gamma\gamma}$. We can also show that the geometric mean is minimized when the ball U_2 is rotated so that the covariance matrix of f is diagonal while the arithmetic mean is unchanged.

5.3. *Higher dimensions*

The general version of Lemma 5.5 is the following.

Lemma 5.9. Suppose $a_1, a_2, \ldots, a_d, \alpha_1, \alpha_2, \ldots, \alpha_d$ are nonnegative reals, and $\alpha_1 + \ldots + \alpha_d \leq 1$. Then

$$(a_1\alpha_1 + \ldots + a_d\alpha_d)^k \leq a_1^k\alpha_1 + \ldots + a_d^k\alpha_d$$

Proof: Prove by induction on k and using the condition $\sum_{i=1}^{d} \alpha_i \leq 1$.

$$\begin{aligned}
&(a_1\alpha_1 + \ldots + a_d\alpha_d)^k \\
&= (a_1\alpha_1 + \ldots + a_d\alpha_d)^{k-1}(a_1\alpha_1 + \ldots + a_d\alpha_d) \\
&\leq \left[a_1^k\alpha_1 + \ldots + a_d^k\alpha_d\right] \\
&\quad - \left(\sum_{i=1}^{d} a_i^k\alpha_i(1 - \alpha_i)\right) + \left(\sum_{i,j} a_i^{k-1} a_j\alpha_i\alpha_j\right) \\
&\leq \left[a_1^k\alpha_1 + \ldots + a_d^k\alpha_d\right] \\
&\quad - \left(\sum_{i=1}^{d} \sum_{j\neq i} a_i^k\alpha_i\alpha_j\right) + \left(\sum_{i,j} a_i^{k-1} a_j\alpha_i\alpha_j\right) \\
&= \left[a_1^k\alpha_1 + \ldots + a_d^k\alpha_d\right] - \\
&\quad \left[\left(\sum_{i,j} \alpha_i\alpha_j(a_i^{k-1} - a_j^{k-1})(a_i - a_j)\right)\right]
\end{aligned}$$

$$\leq a_1^k \alpha_1 + \ldots + a_d^k \alpha_d$$

□

In d dimensions, the second moments of f on U_d with respect to the coordinate axes are given by

$$\mu_{ii}(f) = \frac{1}{|U_{d-1}|} \int_{w \in U_{d-1}} f(w)(i(w))^2 dA$$

where $i(w)$ is the ith coordinate of point w on U_{d-1}. Note that $\sum_{i=1}^{d} i^2(w) = 1$ for all w on U_{d-1}. The expected cost of sampling from the uniform distribution is given by

$$UC(f) = \sum_{i=1}^{d} \mu_{ii}(f)$$

Let $E_{\bar{a}}$ denote the d dimensional ellipsoid defined by

$$E_{\bar{a}} = \{(x_1, \ldots, x_d) \mid \sum_{i=1}^{d} \frac{x_i^2}{a_i^2} \leq 1.\}$$

Similar to Theorem 5.8, we have,

Theorem 5.10. If $a_i = (1/\mu_{ii})^{1/d}$ $(1 \leq i \leq d)$, then the expected cost of the elliptic distribution

$$Expect(E_{\bar{a}}, f) \leq d \left(\prod_{i=1}^{d} \mu_{ii} \right)^{1/d}$$

Moreover, since [4]

$$d \left(\prod_{i=1}^{d} \mu_{ii} \right)^{1/d} \leq \sum_{i=1}^{d} \mu_{ii}$$

we have

$$Expect(E_{\bar{a}}, f) \leq UC(f)$$

6. Great circles and duality

The preceding sections have dealt with using the second moments to guide the solution to the primal problem. For application to geometric separators, we need to deal with the dual

problem, in which the cost of a given point w (on the unit sphere) is not $f(w)$, but instead is the sum (integration) of the points on the great circle normal to w (that is, $\int_{v \in GC(w)} f(v) ds^{d-1}$). For presentation, we focus on the two dimensional problem.

For a density function f on the unit 2-sphere, let $\mathcal{M}(f)$ be the moment matrix (3×3 matrix) of f, defined by

$$\mathcal{M}(f) = \begin{pmatrix} \mu_{\alpha\alpha} & \mu_{\alpha\beta} & \mu_{\alpha\gamma} \\ \mu_{\beta\alpha} & \mu_{\beta\beta} & \mu_{\beta\gamma} \\ \mu_{\gamma\alpha} & \mu_{\gamma\beta} & \mu_{\gamma\gamma} \end{pmatrix}$$

where

$$\mu_{ij} = \frac{1}{|U_2|} \int_{w \in U_2} f(w) i(w) j(w) dA, \; i, j \in \{\alpha, \beta, \gamma\}$$

Let $g_f(w) = 1 / \pi \int_{v \in GC(w)} f(v) ds$ denote the dual cost of f, and $\mathcal{D}(f) = \mathcal{M}(g_f)$ be the moment matrix of the dual cost g_f of f. Let $\mu(f) = UC(f) = \int_{w \in U_2} f(w) dA$ denote the expected cost from sampling f with the uniform distribution on U_2. Let I_d denote the $d \times d$ identity matrix.

The following lemma relates the moment matrix of the primal to that of the dual, and is interesting in its own right.

Lemma 6.1. (Primal & Dual) For all density functions f on U_2, $\mathcal{M}(f) + \mathcal{D}(f) = \mu(f) I_3$.

In the proof Lemma 6.1, we will use the following propositions. These propositions follows directly from the definitions.

Proposition 6.2. (Linearity) Let f_1 and f_2 be two density functions on U_2, then $\mathcal{M}(f_1 + f_2) = \mathcal{M}(f_1) + \mathcal{M}(f_2)$ and $\mathcal{D}(f_1 + f_2) = \mathcal{D}(f_1) + \mathcal{D}(f_2)$.

Proposition 6.3. (Primal Rotation) Let f be a density function on U_2 and R a 3×3 unitary matrix. Let h be the resulting density function when applying R to U_2, i.e., mapping point (α, β, γ) to $(\alpha, \beta, \gamma)R$. Then

$$\mathcal{M}(h) = R^T \mathcal{M}(f) R$$

Proposition 6.4. (Dual Rotation) Let f be a density function on U_2 and R a 3×3 unitary matrix. Let h be the resulting density function when applying R to U_2, i.e., mapping point (α, β, γ) to $(\alpha, \beta, \gamma)R$. Then

$$\mathcal{D}(h) = R^T \mathcal{D}(f) R$$

In order to motivate the proof of Lemma 6.1, we first give a proof to the point-density version of the lemma. The point-density version is the one often used in the practical application of the moments of inertia. Let $P = \{p_1, p_2, \ldots, p_n\}$ be a set of n points of weights $\{w_1, w_2, \ldots, w_n\}$, respectively, on U_2. Assume that $p_i = (x_i, y_i, z_i)^T$. The moment matrix of P is given by

$$\begin{aligned}\mathcal{M}(P) &= \sum_{i=1}^{n} \mathcal{M}(p) \\ &= \sum_{i=1}^{n} w_i(p_i p_i^T) \\ &= \sum_{i=1}^{n} w_i \begin{pmatrix} x_i^2 & x_i y_i & x_i z_i \\ x_i y_y & y_i^2 & y_i z_i \\ x_i z_i & y_i z_i & z_i^2 \end{pmatrix},\end{aligned}$$

where $(p_i p_i^T)$ is the outer-product of the vector p_i.

To define the dual moment matrix, we use the duality between great circles and points on the unit sphere: Each great circle GC can be identified with a pair of points p_{GC} and q_{GC} on U_2 that lie on the normal to GC. (See Figure 2). Simply from the definition,

Proposition 6.5. (Duality) For each pair of great circles $\mathcal{G}$ and $\mathcal{G}'$ of U_d, $\mathcal{G}$ contains $p_{\mathcal{G}'}$ (and hence $q_{\mathcal{G}'}$ as well) if and only if $\mathcal{G}'$ contains $p_{\mathcal{G}}$ (and hence $q_{\mathcal{G}}$).

Let $p = (x, y, z)$ be a point of density w on U_2. By Proposition 6.5, a great circle passes through p if its normal points are on the great circle normal to p. Hence, the density at p defines a density function in the dual space: The dual-density of a point is w if it is on the great circle normal to p; otherwise its dual density is equal to 0. Note that there are infinite number of points with non-zero dual-density. We need to properly define the dual moment matrix based on the concepts of calculus. To motivate, we relate a point-density to a continuous density function. Notice that a point is an object of zero dimension, whereas the sphere U_2 is a two dimensional object. The basic idea to make a point-density con-

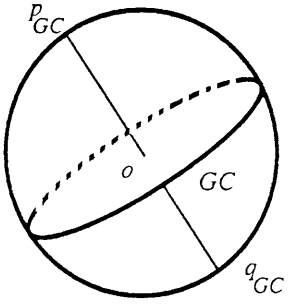

Figure 2. A great circle and its normal points.

tinuous, in the view of calculus, is to treat a point p as an infinitely small patch dA of area ds^2 and the value of the continuous function in dA is $w/(ds^2)$. Let ϕ_p be such a continue function. Then the moment matrix of ϕ_p, $\mathcal{M}(\phi_p)$, is equal to the moment matrix of p, $\mathcal{M}(p)$.

In the dual space, each great circle has dimension one. Based on the same observation as in the paragraph above, we should "give" a great circle a width ds. The dual-density induced by a point-density w at p is then $w/(\pi ds)$ at points on the great circle normal to p. In other words, we uniformly distribute the point density over the points on the great circle normal to it. The term π in the denominator is the perimeter of a great circle. We now calculate the moment matrix of the dual-density.

The simplest case is that of a point density w at $p = (0, 0, 1)$. The primal moment matrix is

$$\mathcal{M}(p) = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & w \end{pmatrix}$$

The dual-density would be $w/(\pi ds)$ at every point on the great circle (of width ds) on the xy-plane (that is, normal to (0, 0, 1)), since the great circle normal to each such point would pass through (0, 0, 1). Let G be the great circle normal to p. The dual moment matrix is

$$\begin{aligned} D(p) &= \int_{q \in G} (w/(\pi ds))(qq^T) ds^2 \\ &= (w/\pi) \int_{q \in G} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} ds \\ &= \begin{pmatrix} w & 0 & 0 \\ 0 & w & 0 \\ 0 & 0 & 0 \end{pmatrix} \\ &= wI_3 - \mathcal{M}(p) \end{aligned}$$

Lemma 6.1 is true in this case for a point density.

Now we look at the discrete case where we have a point-density w at point $p = (\alpha, \beta, \gamma)$ of U_2. The moment matrix of the point density is $w(\alpha, \beta, \gamma)^T (\alpha, \beta, \gamma)$.

Let R be a unitary matrix that transforms the point (0, 0, 1) to (α, β, γ) (rotation or reflection). Then $(\alpha, \beta, \gamma) = (0, 0, 1)R$.

Since R also maps the great circle of (0, 0, 1) to that of (α, β, γ), we have that the dual moment matrix is, by Proposition 6.4, $R^T D(m)R = wI_3 - w(\alpha, \beta, \gamma)^T (\alpha, \beta, \gamma)$. Lemma 6.1 is again true here.

By the linearity of the moment matrix (Proposition 6.2), we have,

$$\begin{aligned}\mathcal{D}(P) &= \sum_{i=1}^{n} \mathcal{D}(p) \\ &= \sum_{i=1}^{n} (w_i I_3 - w_i \mathcal{M}(p_i)) \\ &= WI_w + \mathcal{M}(P),\end{aligned}$$

where $W = \sum_{i=1}^{n} w_i$. Therefore, we have proved the following version of Lemma 6.1 for a collection of point densities on U_2.

Lemma 6.6. (point primal & dual) Let $P = \{p_1, p_2, \ldots, p_n\}$ be a set of n points of weights $\{w_1, w_2, \ldots, w_n\}$, respectively, on U_2. Then $\mathcal{M}(P) + \mathcal{D}(P) = WI_3$, where $W = \Sigma_{i=1}^{n} w_i$.

Lemma 6.6 also has another interesting interpretation, that is, $\mathcal{D}(P)$ is exactly the same as the inertia matrix of P. (See Section 2.2 for the definition of the inertia matrix.)

To prove Lemma 6.1, we discretize U_2 into a collection of infinitely small patches and use summation to approximate the integration. We then reduce the continuous case to the point-density case. Lemma 6.1 shows that the dual-moment matrix is the same as the inertia matrix defined in Section 5.1.

We now have the following lemma, which follows directly from Lemma 6.1.

Lemma 6.7. Let f_1 and f_2 be two density functions on U_2. If $\mathcal{M}(f_1) = \mathcal{M}(f_2)$, then $\mathcal{D}(f_1) = \mathcal{D}(f_2)$.

Following from Theorem 5.8 and Lemma 6.1, we obtain a theorem for the Dual Problem of Section 3.

Theorem 6.8. If

$$a = (1/(\mu - \mu_{\alpha\alpha}))^{1/3}$$

$$b = (1/(\mu - \mu_{\beta\beta}))^{1/3}$$

$$c = (1/(\mu - \mu_{\gamma\gamma}))^{1/3}$$

then the expected dual cost for the elliptic distribution is

$$DEC_{abc} = 3[(\mu - \mu_{\alpha\alpha})(\mu - \mu_{\beta\beta})(\mu - \mu_{\gamma\gamma})]^{1/3},$$

compared with the expected dual cost for the uniform distribution of

$$DUC = [(\mu - \mu_{\alpha\alpha}) + (\mu - \mu_{\beta\beta}) + (\mu - \mu_{\gamma\gamma})] = 2\mu.$$

Therefore, $DEC_{abc} \leq DUC$.

The theorem above can be generalized to higher dimensions.

Theorem 6.9 If $a_i = (1/(\mu^* - \mu_{ii}))^{1/d}$, $(1 \leq i \leq d)$, then the expected dual elliptic cost is $DEC_{\bar{a}} = d\,[\prod_{i=1}^{d} (\mu - \mu_{ii})]^{1/d}$ compared with the expected dual uniform cost of $DUC = [\sum_{i=1}^{d} (\mu - \mu_{ii})] = (d-1)\mu$. Therefore, $DEC_{\bar{a}} \leq DUC$.

7. Applications in Geometric Graph Separators

As shown in [30], all well-shaped finite element meshes and some graphs from computational geometry, such as k-nearest neighbor graphs have a geometric characterization. We now review this characterization and show how to apply second moments to improve the sphere separator algorithm.

7.1. Neighborhood systems and their separator

The characterization is based on the notion of a neighborhood system. A d-dimensional *neighborhood system* $\phi = \{B_1,\ldots, B_n\}$ is a finite collection of balls in $\mathbb{R}^d$. Let p_i be the center of B_i $(1 \leq i \leq n)$ and call $P = \{p_1,\ldots, p_n\}$ the centers of ϕ. For each point $p \in \mathbb{R}^d$, let the ply of p, denoted by $\text{ply}_\phi(p)$, be the number of balls from ϕ that contains p. ϕ is a *k-ply* neighborhood system if for all p, $\text{ply}_\phi(p) \leq k$.

Each $(d-1)$-dimensional sphere S in $\mathbb{R}^d$ partitions ϕ into three subsets: $\phi_I(S)$, $\phi_E(S)$, and $\phi_O(S)$, which are the balls that are in the interior of S, in the exterior of S, and that intersect S, respectively. The cardinality of $\phi_O(S)$ is called the intersection number of S, denoted by $l_\phi(S)$.

Notice that the removal of $\phi_O(S)$ splits ϕ into two subsets: $\phi_I(S)$ and $\phi_E(S)$, such that no ball in $\phi_I(S)$ intersects any ball in $\phi_E(S)$ and vice versa. In analogy to separators in graph theory, $\phi_O(S)$ can be viewed as a separator of ϕ.

Definition 7.1. (Sphere Separators) A $(d-1)$-sphere S is an $f(n)$-**separator** that δ-**splits a neighborhood system** ϕ if $l_\phi(S) \leq f(n)$ and $|\phi_I(S)|, |\phi_E(S)| \leq \delta n$, where f is a positive function and $0 < \delta < 1$.

Simply by definition, each $f(n)$-sphere separator that δ-splits also naturally induces a vertex separator of the intersection graph. So, we will just focus on the sphere separators of neighborhood systems. The following theorem is proved in [30]:

Theorem 7.2. (Sphere Separator Theorem) Suppose $\phi = \{B_1,\ldots, B_n\}$ is a k-ply neighborhood system in $\mathbb{R}^d$. Then there is a $(d-1)$-sphere S intersecting at most $O(k^{1/d} n^{(d-1)/d})$ balls from ϕ such that both $|\phi_I(S)|$ and $|\phi_E(S)|$ are at most $\leq (d+1)/(d+2)\, n$, where $\phi_I(S)$ and $\phi_E(S)$ are those balls that are in the interior and in the exterior of S, respectively.

We now review the construction of [30] to show where the second moments can be used. To find a small cost sphere separator with balanced splitting ratio, we first map the neighborhood system ϕ conformally onto the unit sphere U_d so that each d-dimensional

hyperplane containing the center of U_d δ-splits $\phi = \{B_1,\ldots, B_n\}$ $(d + 1)/(d + 2) < \delta < 1)$. This step can be performed in random constant time using a randomized center point algorithm [3, 36]. Now, each B_i is mapped to a patch D_i on U_d, whose boundary C_i has the shape of a (d-1)-sphere. Each great circle induces a balanced sphere separator and the cost of a great circle is the number of patches it intersects. Therefore, the goal here is to find a great circle with a small intersection number. In the remainder of this section, we will identify B_i with D_i, and assume that $\phi = \{B_1,\ldots, B_n\}$ is given on the unit d-sphere.

7.2. Exact moment by duality

To apply second moments, we use the duality between great circles and points on the unit sphere: Each great circle GC can be identified with a pair of points p_{GC} and q_{GC} on U that lie on the normal to GC. (See Figure 2).

Define a great belt to be the set of points of U_d that lie between a pair of parallel hyperplanes located symmetrically about the center of U_d. The next lemma follows from the definition. See [8, 36] for a proof.

Lemma 7.3. Suppose $\phi = \{B_1,\ldots, B_n\}$ is a neighborhood system on U_d. Then for each $1 \leq i \leq n$, there is a great belt R_i such that a great circle $\mathcal{G}$ intersects B_i iff $p_{\mathcal{G}}$ and $q_{\mathcal{G}}$ is contained in R_i. (See Figure 3).

Therefore, to find a small sphere separator, it is equivalent to find a point in U_d that is covered by a small number of great belts. We can define a density function on the unit sphere U_d for which the density of a point is the number of great belts that cover it. In other words, each point of each great belt has a unit local density and the global density

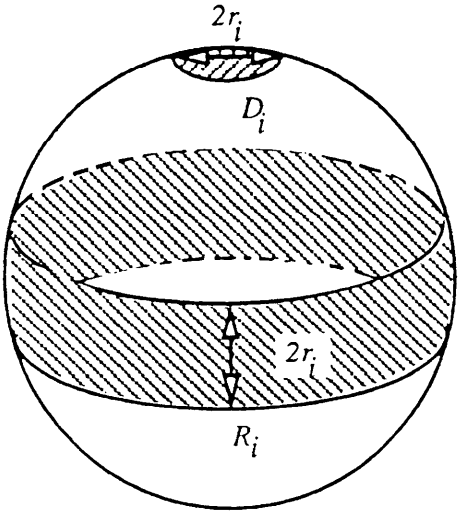

Figure 3. A ball on the sphere and its dual great belt.

is the sum of the local densities. Hence we reduce the sphere separator problem to the primal case of the optimization problem discussed in Section 3. We can first compute the great belts and use them to find the moment matrix. Then the construction and results of Section 5 can be applied directly. We refer to this construction as the *dual-primal* method, because we first use duality to find great belts and then compute the moment matrix.

7.3. Approximation by dual moments and geometric sampling

A more efficient way of applying second moments is to use the primal-dual method, which finds the moment matrix using the primal information, rather than great belts, and then applies the construction of Section 6.

The advantages of this method are twofold: (1) we do not need to compute great belts and (2) we can use geometric sampling to approximate the moment matrix.

We use the following observation made in [30, 36]. Because $\phi = \{B_1,\ldots, B_n\}$ is a k-ply neighborhood system on U_d, then there exists a constant C such that the number of balls from ϕ whose radius is at least γ is bounded from above by Ck/γ^d. Therefore, the radii of most of the balls are small, so that they can be treated as points, given by their centers. Formally, the weight of each point should be the diameter of its corresponding ball. Now, our input becomes a set of points $Q = \{q_1 \ldots q_n\}$ on U_d and weights $w_1 \ldots w_n$. The moment matrix is then $\mathcal{M} = \sum_{i=1}^{n} w_i(q_i q_i^T)$, where $(q_i q_i^T)$ is the outer product of the vector q_i. We can then use geometric sampling to approximate the moment matrix. This leads to the algorithm in Figure 4. Note that in the algorithm we have set the weights to one.

In Figure 4, *stereo-up* is the standard stereographic projection. For more detail on other steps such as finding center points, finding a good conformal map, and finding a vertex separator from a sphere separator, see [30, 31].

Theorem 7.4. The algorithm above finds a sphere separator whose expected cost is less than or equal to that found by the algorithm of [30] which returns a random great circle.

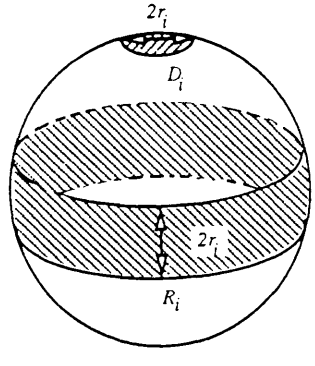

Figure 4. The moment-based sphere separator algorithm.

Table 1. Separator sizes obtained from 7 trials on airfoil1, a 2D mesh with 4253 nodes and 12289 edges.

method	separator size						
Moment	62	60	64	73	69	64	69
Uniform	114	155	145	133	119	278	91

Table 2. Separator sizes obtained from 7 trials on airfoil2, a 2D mesh with 4720 nodes and 13722 edges.

method	separator size						
Moment	99	96	94	94	96	103	93
Uniform	178	188	129	190	154	132	211

Table 3. Separator sizes obtained from 7 trials on airfoil3, a 2D mesh with 15606 nodes and 45878 edges.

method	separator size						
Moment	168	154	176	190	174	174	160
Uniform	174	276	181	191	207	232	219

Table 4. Separator sizes obtained from 7 trials on brack2, a 3D mesh with 62631 nodes and 335240 edges.

method	separator size						
Moment	752	746	719	712	768	715	770
Uniform	1083	2761	5724	1740	4401	1784	2285

8. Experimental results

We implemented the moment-based method by modifying the existing Mat-lab geometric partitioning code of Gilbert and Teng [17] to sample sphere separators from an elliptic distribution based on the second moments. To compare the results of the original and modified algorithms, implementation was performed by angle updating. Hence, in each trial, a random angle was selected from the uniform distribution, and the angle was updated to yield an angle from the elliptic distribution. The parameters used were more aggressive than those predicted by theory.

Our preliminary experiments show that using second moments improves the sphere separator algorithm for well-shaped meshes in both two and three dimensions. In each trial, the result of sampling from an elliptic distribution improved the result of sampling from the uniform distribution.

Tables 1–4 give the sizes of edge-separators found in first seven trials of the sphere separator algorithm when the moment-based distribution (moment method) and the uniform distribution (uniform method) were used, respectively. In all cases, the best separator was found using the moment-based distribution.

Acknowledgments

We would like to thank John Gilbert, Bruce Hendrickson, and Horst Simon for references to the work of Nour-Omid.

References

Alon, N., P. Seymour, and R. Thomas. A separator theorem for non-planar graphs. In *Proceedings of the 22th Annual ACM Symposium on Theory of Computing*, Maryland, May 1990. ACM.

Blelloch, G., A. Feldmann, O. Ghattas, J. Gilbert, G. Miller, D. R. O'Hallaron, E. Schwabe and J. Schewchuk, S.-H. Teng. (1996). Automated parallel solution of unstructured PDE problems. *CACM*, invited submission, to appear.

Clarkson, K., D. Eppstein, G. L. Miller, C. Sturtivant, and S.-H. Teng. Approximating center point with iterated radon points. In *Proceedings of 9th ACM Symposium on Computational Geometry*, pp. 91–98, San Diego, May, 1993.

Crawford, G. E. Elementary proof that the arithmetic mean of any number of positive quantities is greater than the geometric mean. *Proc. Edinburgh Math. Soc.*, 17:2–4, 1899–1900.

Djidjev, H. N. On the problem of partitioning planar graphs. *SIAM J. Alg. Disc. Math.*, 3(2):229–240, June 1982.

Donath, W. E. (1988). Logic partitioning. In B. T. Preas and M. J. Lorenzetti, eds., *Physical Design Automation of VLSI Systems*, pp. 65–86. Benjamin/Cummings.

Donath, W. E., and A. J. Hoffman. (1972). Algorithms for partitioning of graphs and computer logic based on eigenvectors of connection matrices. *IBM Technical Disclosure Bulletin*, 15:938–944.

Eppstein, D., G. L. Miller, and S.-H. Teng. (1993). A deterministic linear time algorithm for geometric separators and its applications. In *Proceedings of 9th ACM Symposium on Computational Geometry*, pp. 99–108, San Diego, May.

Farhat, C., and M. Lesoinne. (1993). Automatic partitioning of unstructured meshes for the parallel solution of problems in computational mechanics. Int. J. Num. Meth. Eng. 36:745–764.

Farhat, C., and H. Simon. (1993). TOP/DOMDEC—a software tool for mesh partitioning and parallel processing. Technical Report, NASA Ames Research Center.

Frieze, A. M., G. L. Miller, and S.-H. Teng. (1992). Separator based parallel divide and conquer in computational geometry. In *4th Annual ACM Symposium on Parallel Algorithms and Architectures*, pp 420–430.

Gazit, H. (1986). An improved algorithm for separating a planar graph. Manuscript, Department of Computer Science, University of Southern California.

Gazit, H., and G. L. Miller. A parallel algorithm for finding a separator in planar graphs. In *28st Annual Symposium on Foundation of Computation Science, IEEE*, 238–248, Los Angeles, October 1987.

George, J. A. (1973). Nested dissection of a regular finite element mesh. *SIAM J. Numerical Analysis*, 10:345–363.

George, A., and J. W. H. Liu. (1981). *Computer Solution of Large Sparse Positive Definite Systems*. Prentice-Hall.

Gilbert, J. R., J. P. Hutchinson, and R. E. Tarjan. (1984). A separation theorem for graphs of bounded genus. *J. Algorithms*, 5:391–407.

Gilbert, J. R., G. L. Miller, and S.-H. Teng. (1995). Geometric mesh partitioning: Implementation and experiments. In *International Conference of Parallel Processing*, pp 418–427.

Heath, M., and P. Raghavan. (1994). A cartesian parallel nested dissection algorithm. To appear in *SIAM Journal on Matrix Analysis and Applications*.

Hendrickson, B., and R. Leland. (1993). Multidimensional spectral load balancing. Technical Report, Sandia National Laboratories, SAND93-0074.

Hendrickson, B., and R. Leland. (1993). A multilevel algorithm for partitioning graphs. Technical Report SAND93-1301, Sandia National Laboratories, Albuquerque, NM.

Hendrickson, B., and R. Leland. (1993). The Chaco user's guide, Version 1.0. Technical Report SAND93-2339, Sandia National Laboratories, Albuquerque, NM.

Jordan, C. (1869). Sur les assemblages de lignes. *Journal Reine Angew. Math*, 70:185–190.

Kernighan B. W., and S. Lin. (1970). An efficient heuristic procedure for partitioning graphs. *Bell Sys. Tech. J.*, 49: 291–307.

Leighton, F. T. (1983). *Complexity Issues in VLSI*. Foundations of Computing. MIT Press, Cambridge, MA.

Leighton, F. T., and S. Rao. (1988). An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. In *29th Annual Symposium on Foundations of Computer Science*, pp 422–431.

Leiserson, C. E. (1983). *Area Efficient VLSI Computation*. Foundations of Computing. MIT Press, Cambridge, MA.

Lipton, R. J., D. J. Rose, and R. E. Tarjan. (1979). "Generalized nested dissection". *SIAM J. on Numerical Analysis*, 16:346–358.

Lipton, R. J., and R. E. Tarjan. (1979). "A separator theorem for planar graphs". *SIAM J. of Appl. Math.*, 36(April), 177–189.

Miller, G. L. (1986). Finding small simple cycle separators for 2-connected planar graphs. *Journal of Computer and System Sciences*, 32(3)(June), 265–279.

Miller, G. L., S.-H. Teng, W. Thurston, and S. A. Vavasis. (1996). "Automatic Mesh Partitioning." In A. George, J. Gilbert, and J. Liu, editors, *Sparse Matrix Computations: Graph Theory Issues and Algorithms*, IMA Volumes in Mathematics and its Applications. Springer-Verlag.

Miller, G. L., S.-H. Teng, W. Thurston, and S. A. Vavasis. (1996). Finite element meshes and geometric separators. *SIAM J. Scientific Computing*, to appear.

Nour-Omid, B., A. Raefsky, and G. Lyzenga. (1987). Solving finite element equations on concurrent computers. in A. K. Noor, ed., *Parallel Computations and Their Impact on Mechanics*, The American Society of Mechanical Engineers, AMD-Vol. 86, 209–228.

Pan, V., and J. Reif. (1985). Efficient parallel solution of linear systems. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*, pages 143–152, Providence, RI, May. ACM.

Pothen, A., H. D. Simon, and K.-P. Liou. (1990). Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Mat. Anal. Appl.*, 11(3):430–452.

Simon, H. D. (1991). Partitioning of unstructured problems for parallel processing. *Computing Systems in Engineering* 2(2/3):135–148.

Teng, S.-H. (1991). *Points, Spheres, and Separators: a unified geometric approach to graph partitioning*. Ph.D. Thesis, Carnegie Mellon University, CMU-CS-91-184.

Ungar, P. (1951). A theorem on planar graphs. *Journal London Math Soc.* 26:256–262.

Williams, R. D. (1991). Performance of dynamic load balancing algorithms for unstructured mesh calculations *Concurrency: Practice and Experience*, 3(5):457–481.