

Geometric Separators and the Parabolic Lift

Donald R. Sheehy*

Abstract

A geometric separator for a set U of n geometric objects (usually balls) is a small (sublinear in n) subset whose removal disconnects the intersection graph of U into roughly equal sized parts. These separators provide a natural way to do divide and conquer in geometric settings. A particularly nice geometric separator algorithm originally introduced by Miller and Thurston has three steps: compute a centerpoint in a space of one dimension higher than the input, compute a conformal transformation that “centers” the centerpoint, and finally, use the computed transformation to sample a sphere in the original space. The output separator is the subset of S intersecting this sphere. It is both simple and elegant. We show that a change of perspective (literally) can make this algorithm even simpler by eliminating the entire middle step. By computing the centerpoint of the points lifted onto a paraboloid rather than using the stereographic map as in the original method, one can sample the desired sphere directly, without computing the conformal transformation.

1 Geometric Separators

A spherical geometric separator of a collection of n balls in \mathbb{R}^d is a sphere S that has at least $\frac{n}{d+2}$ balls centered inside, at least $\frac{n}{d+2}$ centered outside, and intersects at most $O(n^{1-\frac{1}{d}})$ of them (see Section 2 for a formal definition). The existence of such separators in two and three dimensions was established by Miller and Thurston, though their method was quickly adapted to higher dimensions. Across a series of papers, Miller, Thurston, Teng, and Vavasis laid out the theory of geometric separators and their applications to scientific computing [14, 15, 13, 10, 12]. This line of work is a treasure trove for computational geometers as it hinges on a novel trick that combines projective and combinatorial geometry to solve an important algorithmic problem, solving linear systems arising in finite element analysis.

More generally, geometric separators give a natural way to do divide and conquer for geometric problems. They have been applied to various nearest neighbor search problems [11] as well as to mesh compression [1]. Other variations of geometric separators have been used

for the Traveling Salesman and Minimum Steiner Tree problems in geometric settings [18], or for packing and piercing problems [2].

The Miller-Thurston algorithm for computing a geometric separator maps the n points (the centers of the balls) to a unit d -sphere in \mathbb{R}^{d+1} via a stereographic map. It then computes a centerpoint, which is a geometric generalization of a median (a formal definition is given in Section 2). There exists a conformal transformation of the points in \mathbb{R}^{d+1} so that this centerpoint will lie exactly at the origin. To output a separator, one samples a random unit vector in \mathbb{R}^{d+1} . The hyperplane through the origin normal to this vector intersects the unit d -sphere at a $(d-1)$ -sphere. The output is just the stereographic projection of this $(d-1)$ -sphere back to \mathbb{R}^d . With high probability, such a sphere will be a geometric separator.

The one aspect of this algorithm that was left to the reader, was the linear algebra required to compute the desired conformal transformation. In the original paper, it was simply asserted that it exists. Later papers explained that it can be computed via Householder transformations and cited a textbook on matrix computations. In this paper, we show that this phase of the algorithm is entirely unnecessary. By working initially with the parabolic lifting

$$\mathbf{p} \mapsto \left[\frac{\mathbf{p}}{\|\mathbf{p}\|^2} \right],$$

rather than the stereographic map, the desired sphere can be sampled directly.

Related work In addition to the previously mentioned work on sphere separators and their applications by various combinations of Miller, Thurston, Teng, and Vavasis, the generalization to other types of contact graphs and other shapes of separators (in particular hypercubes) was performed by Smith and Wormald [18]. This method was refined slightly by Chan to apply separators to geometric hitting set problems [2]. Eppstein et al. gave a linear-time, deterministic algorithm for finding geometric separators based on core sets [5]. Har-Peled gave a simple proof of the existence of geometric separators for interior-disjoint disks in the plane (that easily extends to higher dimensions) [6].

*Department of Computer Science and Engineering, University of Connecticut, don.r.sheehy@gmail.com

2 Definitions

Points We treat points in Euclidean space as column vectors. As the algebra will be in both \mathbb{R}^d as well as \mathbb{R}^{d+1} , we adopt the convention that a boldface vector \mathbf{p} is a vector of \mathbb{R}^d and an overline such as in $\overline{\mathbf{p}}$ indicates a vector in \mathbb{R}^{d+1} . In particular, we write $\mathbf{0}$ to indicate the zero vector in \mathbb{R}^d . Scalars are italic and when it is useful, we add a subscript $d + 1$ as in p_{d+1} to indicate a scalar that is the $(d + 1)$ st coordinate of a vector in \mathbb{R}^{d+1} . So, for example, we will write $\overline{\mathbf{p}} = \begin{bmatrix} \mathbf{p} \\ p_{d+1} \end{bmatrix}$ to indicate that $\overline{\mathbf{p}} \in \mathbb{R}^{d+1}$ with its first d coordinates matching those of \mathbf{p} and last coordinate equal to p_{d+1} . Whenever we speak of \mathbb{R}^d as a subspace of \mathbb{R}^{d+1} , it is always assumed that we mean the hyperplane $\{\overline{\mathbf{p}} \mid p_{d+1} = 0\}$ in \mathbb{R}^{d+1} .

Projections Given a point $\overline{\mathbf{f}} = \begin{bmatrix} \mathbf{f} \\ f_{d+1} \end{bmatrix}$, we define $\Pi_{\overline{\mathbf{f}}}^1$ to be the stereographic map from \mathbb{R}^d to the sphere centered at $\overline{\mathbf{f}}$ with radius f_{d+1} . That is, $\Pi_{\overline{\mathbf{f}}}^1(\mathbf{p})$ is the intersection of the sphere with the line through $\begin{bmatrix} \mathbf{p} \\ 0 \end{bmatrix}$ and the north pole, $\begin{bmatrix} \mathbf{f} \\ 2f_{d+1} \end{bmatrix}$ (other than the pole itself). Similarly, we define $\Pi_{\overline{\mathbf{f}}}^\infty$ to be the parabolic lifting map from \mathbb{R}^d to \mathbb{R}^{d+1} that lifts \mathbf{p} to $\overline{\mathbf{p}} = \begin{bmatrix} \mathbf{p} \\ p_{d+1} \end{bmatrix}$ so that $\overline{\mathbf{p}}$ lies on the paraboloid with focal point $\overline{\mathbf{f}}$ and directrix $\{p_{d+1} = -f_{d+1}\}$. The reason for the notation comes from the fact that the parabola is the limiting case of an ellipse formed by moving one focal point to infinity. We will consider more general stereographic projections in Section 4. In all cases, these maps are invertible (except at the north pole).

We abuse notation slightly and let $\Pi_{\overline{\mathbf{f}}}^1(S)$ denote the set $\{\Pi_{\overline{\mathbf{f}}}^1(\mathbf{p}) \mid \mathbf{p} \in S\}$ and similarly for $\Pi_{\overline{\mathbf{f}}}^\infty$. In particular, $\Pi_{\overline{\mathbf{f}}}^1(\mathbb{R}^d)$ denotes the sphere centered at $\overline{\mathbf{f}}$ with radius f_{d+1} and $\Pi_{\overline{\mathbf{f}}}^\infty(\mathbb{R}^d)$ denotes the paraboloid with focal point $\overline{\mathbf{f}}$ and directrix $\{p_{d+1} = -f_{d+1}\}$.

Centerpoints Given n points in \mathbb{R}^d , a centerpoint is a point $\mathbf{c} \in \mathbb{R}^d$ such that any closed halfspace containing more than $\frac{nd}{d+1}$ points also contains \mathbf{c} . The existence of centerpoints follows from Helly’s Theorem and the pigeonhole principle. Let $\text{CENTERPOINT}(P)$ denote the set of all centerpoints of the set P .

It is not known how to find a centerpoint deterministically in polynomial time for point sets in \mathbb{R}^d . However, there is an efficient randomized algorithm [3] known for at least twenty years and some recent work on deterministic approximation algorithms [9, 16].

Sphere Separators A **graph separator** is a subset of vertices in a graph whose removal disconnects the graph. Usually, the goal is to find a small separator that separates the graph into roughly equal sized pieces. The most famous example of graphs admitting small separators is the Planar Separator Theorem of Lipton

and Tarjan [8], which states that a separator of size $O(\sqrt{n})$ is always possible for planar graphs that has at least $\frac{n}{3}$ vertices in each of the resulting components.

The early work on separators was directed at solving linear systems by generalized nested dissection, a method for ordering pivots in Cholesky decomposition of sparse, symmetric, positive definite matrices [7]. Of particular interest were those linear systems arising in the finite element method. These systems reflected the underlying geometric structure of the problem domain and thus it was natural to look to the geometry to find small separators [12]. In particular, it often sufficed to consider various definitions of intersection graphs of systems of balls.

Let $\mathcal{B} = \{B_1, \dots, B_n\}$ be a collection of interior disjoint balls in \mathbb{R}^d . Let S be a sphere in \mathbb{R}^d and let $\mathcal{B}_I(S)$, $\mathcal{B}_E(S)$, and $\mathcal{B}_O(S)$ be the subsets of \mathcal{B} that are interior to, exterior to, and intersecting S respectively. The following theorem is a combination of the main existence result for geometric separators with the algorithm used to prove existence. We state it this way, to make clear that the geometric challenge for computing a separator this way lies in finding a stereographic projection to a sphere that has a centerpoint at the center of the sphere.

Theorem 1 (Sphere Separator Thm.[14, 11, 12])

Let \mathcal{B} be a collection of n interior disjoint balls with centers P . Let $\overline{\mathbf{v}} \in \mathbb{R}^{d+1}$ be chosen uniformly from the unit d -sphere. If $\overline{\mathbf{f}}$ is a point in \mathbb{R}^{d+1} such that $\overline{\mathbf{f}}$ is a centerpoint of $\Pi_{\overline{\mathbf{f}}}^1(P)$ and $H = \{\overline{\mathbf{p}} \mid \overline{\mathbf{v}}^\top(\overline{\mathbf{p}} - \overline{\mathbf{f}}) = 0\}$ is the hyperplane through $\overline{\mathbf{f}}$ normal to $\overline{\mathbf{v}}$, then the sphere

$$S = (\Pi_{\overline{\mathbf{f}}}^1)^{-1}(\Pi_{\overline{\mathbf{f}}}^1(\mathbb{R}^d) \cap H)$$

has the property that

$$|\mathcal{B}_O(S)| = O(n^{1-1/d}), \text{ and}$$

$$|\mathcal{B}_I(S)|, |\mathcal{B}_E(S)| \leq \frac{(d+1)n}{d+2}$$

with probability at least $\frac{1}{2}$.

For ease of exposition, we only give this simplest version of the theorem. However, it has also been proven for k -ply neighborhood systems where the balls are permitted to have up to k -wise interior intersections [11] as well as for α -overlap graphs where two balls are considered neighbors if for both balls increasing the radius of one by a factor of α causes them to intersect [12]. In these cases, the bounds depend on k and α respectively, but the algorithm is the same and so the results of this paper apply to these versions of the theorem as well. The version stated above, when combined with the Koebe-Andreev-Thurston embedding theorem for planar graphs is strong enough to prove the Planar Separator Theorem.

3 The Algorithm

When Archimedes quipped that he could move the earth if given a sufficiently long lever, he transferred the majority of the work to the lever builders of his day. We will do likewise and assume that we are given a long lever in the form of an algorithm to efficiently compute a centerpoint of n points in \mathbb{R}^{d+1} . Given such an algorithm, the heavy lifting will be quite easy, and as with Archimedes, that is entirely the point.

Let $P = \{\mathbf{p}_1, \dots, \mathbf{p}_n\} \subset \mathbb{R}^d$ be a set of points. We first present the Miller-Thurston algorithm for computing a separator and then present the new simplified version.

3.1 The Miller-Thurston Algorithm

Let Π be the stereographic map from \mathbb{R}^d to the unit d -sphere centered at the origin in \mathbb{R}^{d+1} . First, compute

$$\bar{\mathbf{c}} \in \text{CENTERPOINT}(\Pi(P)).$$

Find an orthogonal transformation Q such that

$$Q(\bar{\mathbf{c}}) = \begin{bmatrix} \mathbf{0} \\ \theta \end{bmatrix} \text{ for some } \theta \in \mathbb{R}.$$

Let $D = \sqrt{\frac{1-\theta}{1+\theta}} I$, where I is the identity on \mathbb{R}^d .

Choose a random unit vector $\bar{\mathbf{v}} \in \mathbb{R}^{d+1}$ and let S_0 be the d -sphere formed by intersecting the hyperplane $\{\bar{\mathbf{p}} \mid \bar{\mathbf{v}}^\top \bar{\mathbf{p}} = 0\}$ with the unit d -sphere centered at the origin. **Output** $S = \Pi^{-1}(Q^{-1}\Pi(D\Pi^{-1}(S_0)))$.

3.2 A Simpler Algorithm Using the Parabolic Lift

First, compute

$$\bar{\mathbf{c}} = \begin{bmatrix} \mathbf{c} \\ c_{d+1} \end{bmatrix} \in \text{CENTERPOINT}(\left[\begin{bmatrix} \mathbf{p}_1 \\ \|\mathbf{p}_1\|^2 \end{bmatrix}, \dots, \begin{bmatrix} \mathbf{p}_n \\ \|\mathbf{p}_n\|^2 \end{bmatrix} \right]).$$

Next, choose a random unit vector $\bar{\mathbf{v}} = \begin{bmatrix} \mathbf{v} \\ v_{d+1} \end{bmatrix} \in \mathbb{R}^{d+1}$. Let

$$r = \frac{\sqrt{c_{d+1} - \|\mathbf{c}\|^2}}{|v_{d+1}|}.$$

Output the sphere S with center $(\mathbf{c} - r\mathbf{v})$ and radius r . In the improbable case that $v_{d+1} = 0$, the output is just the hyperplane $\{\mathbf{p} \mid \mathbf{v}^\top(\mathbf{p} - \mathbf{c}) = 0\}$.

3.3 Correctness of the Algorithm

The remainder of this section will prove that the algorithm works. According to the results of Miller et al. [12], all we need is a stereographic projection of \mathbb{R}^d to a d -sphere such that the projected points have a centerpoint at the center of the sphere. The trick presented here is that we will instead show how to find a parabolic lifting that has a centerpoint at the focal point of the paraboloid. Then, we show that if we have a point $\bar{\mathbf{f}}$ such that $\Pi_{\bar{\mathbf{f}}}^\infty(P)$ has a centerpoint at $\bar{\mathbf{f}}$, then $\Pi_{\bar{\mathbf{f}}}^1(P)$

also has a centerpoint at $\bar{\mathbf{f}}$, thus giving the desired map. For any vector $\bar{\mathbf{v}} \in \mathbb{R}^{d+1}$ and any $\mathbf{p} \in \mathbb{R}^d$,

$$\bar{\mathbf{v}}^\top (\Pi_{\bar{\mathbf{f}}}^1(\mathbf{p}) - \bar{\mathbf{f}}) = 0 \text{ iff } \bar{\mathbf{v}}^\top (\Pi_{\bar{\mathbf{f}}}^\infty(\mathbf{p}) - \bar{\mathbf{f}}) = 0.$$

That is, for sampling spheres, there is no difference between using the stereographic map or the parabolic lifting. In fact, we prove a much more general statement about the equivalence of various stereographic projections in Theorem 3.

Theorem 2 *Let \mathcal{B} be a collection of n interior disjoint balls with centers $P = \{\mathbf{p}_1, \dots, \mathbf{p}_n\} \subset \mathbb{R}^d$ and let S be the sphere output by the algorithm in Section 3.2. Then,*

$$|\mathcal{B}_O(S)| = O(n^{1-1/d}), \text{ and} \\ |\mathcal{B}_I(S)|, |\mathcal{B}_E(S)| \leq \frac{(d+1)n}{d+2}$$

with probability at least $\frac{1}{2}$.

Proof. Using Theorem 1, it will suffice to show there exists $\bar{\mathbf{f}}$ such that $S = (\Pi_{\bar{\mathbf{f}}}^1)^{-1}(\Pi_{\bar{\mathbf{f}}}^1(\mathbb{R}^d) \cap H)$ where

$$H = \{\bar{\mathbf{p}} \mid \bar{\mathbf{v}}^\top(\bar{\mathbf{p}} - \bar{\mathbf{f}}) = 0\}$$

and $\bar{\mathbf{f}} \in \text{CENTERPOINT}(\Pi_{\bar{\mathbf{f}}}^1(P))$. The equivalence of different stereographic projections proven in Theorem 3 implies that it will suffice to prove the same facts replacing the stereographic map $\Pi_{\bar{\mathbf{f}}}^1$ with the parabolic lift $\Pi_{\bar{\mathbf{f}}}^\infty$. We will show that the focal point that makes this true is

$$\bar{\mathbf{f}} = \left[\frac{1}{2} \sqrt{c_{d+1} - \|\mathbf{c}\|^2} \right],$$

where $\bar{\mathbf{c}} = \begin{bmatrix} \mathbf{c} \\ c_{d+1} \end{bmatrix}$ is the centerpoint of $\Pi_{\bar{\mathbf{f}}}^\infty(P)$ computed in the first step.

First, we show that $S = (\Pi_{\bar{\mathbf{f}}}^\infty)^{-1}(\Pi_{\bar{\mathbf{f}}}^\infty(\mathbb{R}^d) \cap H)$. We will assume without loss of generality that $v_{d+1} > 0$ since $\begin{bmatrix} \mathbf{v} \\ v_{d+1} \end{bmatrix}$ and $\begin{bmatrix} -\mathbf{v} \\ -v_{d+1} \end{bmatrix}$ are sampled with equal probability and both yield the same sphere. We need only check that S is the orthogonal projection of $\Pi_{\bar{\mathbf{f}}}^\infty(\mathbb{R}^d) \cap H$ to \mathbb{R}^d . The equation for $\Pi_{\bar{\mathbf{f}}}^\infty(\mathbb{R}^d)$ is

$$p_{d+1} = \frac{\|\mathbf{p} - \mathbf{c}\|^2}{2\sqrt{c_{d+1} - \|\mathbf{c}\|^2}} = \frac{\|\mathbf{p} - \mathbf{c}\|^2}{2rv_{d+1}},$$

and the equation for H can be rewritten as

$$p_{d+1} = \frac{\mathbf{v}^\top(\mathbf{c} - \mathbf{p})}{v_{d+1}} + \frac{1}{2} \sqrt{c_{d+1} - \|\mathbf{c}\|^2} \\ = \frac{\mathbf{v}^\top(\mathbf{c} - \mathbf{p})}{v_{d+1}} + \frac{1}{2} rv_{d+1}.$$

So, for a point $\bar{\mathbf{p}} = \begin{bmatrix} \mathbf{p} \\ p_{d+1} \end{bmatrix}$ in the intersection,

$$\frac{\|\mathbf{p} - \mathbf{c}\|^2}{2rv_{d+1}} = \frac{\mathbf{v}^\top(\mathbf{c} - \mathbf{p})}{v_{d+1}} + \frac{1}{2} rv_{d+1}.$$

Multiplying by -2 and completing the square twice gives

$$\|\mathbf{p} - (\mathbf{c} - r\mathbf{v})\|^2 = r^2 v_{d+1}^2 + \|r\mathbf{v}\|^2 = r^2,$$

which is precisely the equation for S .

Now, we show that the focal point $\bar{\mathbf{f}}$ is a centerpoint of $\Pi_{\bar{\mathbf{f}}}^\infty(P)$. There are several different ways to show this, but one nice approach is to observe that for any hyperplane normal to $\bar{\mathbf{v}}$ passing through $\bar{\mathbf{f}}$ there is a corresponding hyperplane H normal to $\begin{bmatrix} \mathbf{c} - r\mathbf{v} \\ -\frac{1}{2} \end{bmatrix}$ passing through $\bar{\mathbf{c}}$. The hyperplane H intersects the paraboloid $\Phi = \{ \begin{bmatrix} \mathbf{p} \\ p_{d+1} \end{bmatrix} \mid p_{d+1} = \|\mathbf{p}\|^2 \}$ in an ellipse that projects orthogonally to the very same sphere of radius r centered at $\mathbf{c} - r\mathbf{v}$. It will suffice to show for any $\bar{\mathbf{p}} = \begin{bmatrix} \mathbf{p} \\ \|\mathbf{p}\|^2 \end{bmatrix} \in H \cap \Phi$, that $\|\bar{\mathbf{p}} - (\mathbf{c} - r\mathbf{v})\|^2 = r^2$. By the definition of H ,

$$(\mathbf{c} - r\mathbf{v})^\top \bar{\mathbf{p}} - \frac{1}{2} \|\bar{\mathbf{p}}\|^2 = (\mathbf{c} - r\mathbf{v})^\top \mathbf{c} - \frac{1}{2} c_{d+1}.$$

Multiplying by -2 and completing the square yields

$$\begin{aligned} \|p - (\mathbf{c} - r\mathbf{v})\|^2 &= \|\mathbf{c} - r\mathbf{v}\|^2 - 2(\mathbf{c} - r\mathbf{v})^\top \mathbf{c} + c_{d+1} \\ &= r^2 \|v\|^2 - \|\mathbf{c}\|^2 + c_{d+1} \\ &= r^2. \end{aligned}$$

The last equality follows from the definition of r and the fact that $\|\bar{\mathbf{v}}\|^2 = 1$.

This implies that $\bar{\mathbf{f}}$ is a centerpoint of the points $\Pi_{\bar{\mathbf{f}}}^\infty(P)$ because every hyperplane passing through $\bar{\mathbf{f}}$ separates the same set of points as the corresponding hyperplane through $\bar{\mathbf{c}}$, which is a centerpoint by definition. \square




Figure 1: Three examples of E_R are illustrated for $R = 1$, $R = 2$, and $R = \infty$ from left to right. In each case, the stereographic map of a single point is illustrated. For the last case, the pole is at infinity, so the result is a vertical projection.

4 Equivalence of Stereographic Maps

Let $\bar{\mathbf{f}} = \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}$. Let R be a real number and consider the ellipsoid $E_R \subset \mathbb{R}^{d+1}$ defined as the points $\bar{\mathbf{p}} = \begin{bmatrix} \mathbf{p} \\ p_{d+1} \end{bmatrix}$ such that

$$\frac{\|\mathbf{p}\|^2}{2R-1} + \frac{(p_{d+1} - R)^2}{R^2} = 1.$$

It is tangent to $\begin{bmatrix} \mathbb{R}^d \\ 0 \end{bmatrix}$ at the origin, has major radius R and all minor radii all equal to $\sqrt{2R-1}$. The stereographic map $\Pi_{\bar{\mathbf{f}}}^R$ through $\begin{bmatrix} \mathbf{0} \\ 2R \end{bmatrix}$, the north pole of this ellipse, from \mathbb{R}^d to E_R is well-defined as is the inverse map (except at $\begin{bmatrix} \mathbf{0} \\ 2R \end{bmatrix}$).

As R goes to infinity, E_R converges to the paraboloid $p_{d+1} = \|\mathbf{p}\|^2/4$. This is perhaps easier to see when one observes that E_R is the set of points equidistant from $\bar{\mathbf{f}}$ and the sphere centered at $\begin{bmatrix} \mathbf{0} \\ 2R-1 \end{bmatrix}$ with radius $2R$. As R goes to infinity, this sphere becomes a plane and the paraboloid is the set of points equidistant from a point and a plane. See Figure 1.

If we intersect the ellipsoid E_R with a hyperplane through one of its focal points and then stereographically project that intersection to \mathbb{R}^d from the pole of the ellipse, then the result is a sphere. The following theorem says that for a fixed hyperplane, it doesn't matter which ellipsoid E_R we started with, we always get the same sphere as illustrated in Figure 2.




Figure 2: The stereographic projection of the intersection of the ellipse and a plane through the focal point is the same as the second focal point is moved up to infinity.

Theorem 3 Let α and β be real numbers in $[1, \infty]$ and let $\bar{\mathbf{f}} \in \mathbb{R}^{d+1}$ be any point. If H is a hyperplane containing $\bar{\mathbf{f}}$, then

$$(\Pi_{\bar{\mathbf{f}}}^\alpha)^{-1}(\Pi_{\bar{\mathbf{f}}}^\alpha(\mathbb{R}^d) \cap H) = (\Pi_{\bar{\mathbf{f}}}^\beta)^{-1}(\Pi_{\bar{\mathbf{f}}}^\beta(\mathbb{R}^d) \cap H).$$

Proof. The sphere $S^\alpha = (\Pi_{\bar{\mathbf{f}}}^\alpha)^{-1}(\Pi_{\bar{\mathbf{f}}}^\alpha(\mathbb{R}^d) \cap H)$ can also be written as

$$S^\alpha = \{ \mathbf{p} \in \mathbb{R}^d \mid \bar{\mathbf{v}}^\top (\Pi_{\bar{\mathbf{f}}}^\alpha(\mathbf{p}) - \bar{\mathbf{f}}) = 0 \},$$

where $\bar{\mathbf{v}}$ is the normal vector of H . By translating the points in \mathbb{R}^d and scaling the space uniformly, we may assume that $\bar{\mathbf{f}} = \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}$. Let R be any real number in the range $[1, \infty]$ and let $\mathbf{q} \in \mathbb{R}^d$ be any point. We will show that $\Pi_{\bar{\mathbf{f}}}^R(\mathbf{q})$ lies on a line through $\bar{\mathbf{f}}$ that does not

depend on R . Thus, $\bar{\mathbf{v}}^\top(\Pi_{\bar{\mathbf{f}}}^\alpha(\mathbf{q}) - \bar{\mathbf{f}}) = 0$ if and only if $\bar{\mathbf{v}}^\top(\Pi_{\bar{\mathbf{f}}}^\beta(\mathbf{q}) - \bar{\mathbf{f}}) = 0$ and therefore, the spheres S^α and S^β must be equal for all values of α and β .

Let $\bar{\mathbf{p}} = \begin{bmatrix} \mathbf{p} \\ p_{d+1} \end{bmatrix}$ be the projection $\Pi_{\bar{\mathbf{f}}}^R(\mathbf{q})$. All of the relevant points $\mathbf{0}$, $\bar{\mathbf{f}}$, \mathbf{p} , and \mathbf{q} are contained in a plane, so we proceed in two dimensions, letting $x = \|\mathbf{p}\|$, $y = p_{d+1}$, and $a = \|q\|$.

Since $\bar{\mathbf{p}}$ lies on E_R , we have $\frac{x^2}{2R-1} + \frac{(y-R)^2}{R^2} = 1$, or equivalently,

$$x^2R^2 + y(y - 2R)(2R - 1) = 0. \quad (1)$$

Since $\bar{\mathbf{p}}$ is also on the line from $\begin{bmatrix} \mathbf{0} \\ 2R \end{bmatrix}$, the north pole of E_R , to $\begin{bmatrix} \mathbf{q} \\ 0 \end{bmatrix}$, its planar coordinates x and y satisfy the equation

$$y = \frac{-2R}{a}x + 2R.$$

It follows that

$$R = \frac{ay}{2(a-x)}, \quad 2R - 1 = \frac{a(y-1) + x}{a-x},$$

$$\text{and } y - 2R = \frac{-xy}{a-x}$$

Plugging these values into (1) gives the following.

$$\frac{x^2(ay)^2}{4(a-x)^2} + \frac{y(-xy)(a(y-1) + x)}{(a-x)^2} = 0.$$

Multiplying by $\frac{4(a-x)^2}{xy^2}$ and collecting terms yields the line

$$(a^2 - 4)x - 4a(y - 1) = 0.$$

We observe that this is the equation of a line through $\bar{\mathbf{f}} = (0, 1)$ as desired. \square

5 Concluding Remarks

The main story of this paper is not entirely new to computational geometry. When the parabolic lifting map was first introduced in the problem of computing Voronoi diagrams by Edelsbrunner and Seidel [4], they replaced previous methods based on the stereographic map. Today, the parabolic lifting is the preferred method for reducing Delaunay triangulations to convex hulls in one higher dimension.

Throughout, we have worked directly with the Euclidean coordinates. This gave the benefit of making the computations more immediately clear, but came at the cost of considering several special cases at infinity. An alternative approach would be to exploit the power of projective geometry, which has been shown to be the natural language for stereographic projections. Even the parabola is just an ellipse in the projective plane. For more on projective geometry, I highly recommend the book by Richter-Gebert [17].

I would like to thank Marc Glisse for helpful conversations and advice on high school algebra.

References

- [1] D. K. Blandford, G. E. Blelloch, D. E. Cardoze, and C. Kadow. Compact representations of simplicial meshes in two and three dimensions. *Int. J. Comput. Geometry Appl.*, 15(1):3–24, 2005.
- [2] T. M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *J. Algorithms*, 46(2):178–189, 2003.
- [3] K. Clarkson, D. Eppstein, G. Miller, C. Sturttivant, and S.-H. Teng. Approximating center points with iterated Radon points. *International Journal of Computational Geometry and Applications*, 6(3):357–377, Sep 1996. invited submission.
- [4] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete & Computational Geometry*, 1(1):25–44, 1986.
- [5] D. Eppstein, G. L. Miller, and S.-H. Teng. A deterministic linear time algorithm for geometric separators and its applications. *Fundamenta Informatica*, 22(4):309–329, 1995.
- [6] S. Har-Peled. A simple proof of the existence of a planar separator, July 10 2011.
- [7] R. J. Lipton, D. J. Rose, and R. E. Tarjan. Generalized nested dissection. *SIAM Journal on Numerical Analysis*, 16(2):346–358, 1979.
- [8] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM J. Appl. Math.*, 36:177–189, 1979.
- [9] G. L. Miller and D. R. Sheehy. Approximate center-points with proofs. *Computational Geometry: Theory and Applications*, 43(8):647–654, 2010.
- [10] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. Automatic mesh partitioning. In A. George, J. Gilbert, and J. Liu, editors, *Graphs Theory and Sparse Matrix Computation*, The IMA Volumes in Mathematics and its Application, pages 57–84. Springer-Verlag, 1993. Vol 56.
- [11] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. Separators for sphere-packings and nearest neighborhood graphs. *Journal of the ACM*, 44(1):1–29, 1997.
- [12] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. Geometric separators for finite-element meshes. *SIAM J. Comput.*, 19(2):364–386, 1998.
- [13] G. L. Miller, S.-H. Teng, and S. A. Vavasis. A unified geometric approach to graph separators. In *Proceedings of the 32nd Annual IEEE Symposium on Foundations of Computer Science*, pages 538–547, 1991.
- [14] G. L. Miller and W. Thurston. Separators in two and three dimensions. In *Proceedings of the 22th Annual ACM Symposium on Theory of Computing*, pages 300–309. ACM, 1990.
- [15] G. L. Miller and S. A. Vavasis. Density graphs and separators. In *Proceedings of the Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 331–336. ACM-SIAM, 1991.

- [16] W. Mulzer and D. Werner. Approximating Tverberg points in linear time for any fixed dimension. In *Proceedings of the 28th ACM Symposium on Computational Geometry*, pages 303–310, 2012.
- [17] J. Richter-Gebert. *Perspectives on Projective Geometry*. Springer, 2011.
- [18] W. D. Smith and N. C. Wormald. Geometric separator theorems and applications. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science*, pages 232–243, 1998.