Median Separators in d Dimensions

J. Sipelstein S.W. Smith J.D. Tygar

December 16, 1988

CMU-CS-88-206

Abstract

Given n sites located in d dimensional Euclidean space, we give an algorithm for processing queries in which for a d dimensional vector we find a $d-1$ dimensional hyperplane L normal to the vector such that an equal number of sites lie on either side of L. Our algorithm uses time $O(\min(n, S(n^d, d-1)))$. and requires preprocessing time $O(n^d \log n)$. Here $S(x, y)$ is the time required to perform deterministic geometric search; that is, given a y dimensional Euclidean region divided into x subregions by y dimensional simplices, the time required to find the subregion in which a given point lies. For two and three dimensions, this gives us $O(\log n)$ algorithms, which is optimal. Our algorithm extends to give us a solution to the k-set problem in an arbitrary number of dimensions.

J. Sipelstein was supported by a National Science Foundation Graduate Fellowship. S. W. Smith was supported by an Office of Naval Research Graduate Fellowship. J. D. Tygar was supported by a National Science Foundation Presidential Young Investigator grant CCR-8858087.

Median Separators in d Dimensions

J. Sipelstein S.W. Smith J. D. Tygar

December 16, 1988

1 Introduction

Consider the *median separator problem*: we have n sites located in d dimensional Euclidean space. These sites are presumed to be in *general position*, i.e., we assume no $j+1$ sites lie in a j-dimensional subspace for any $j < n$ — hence no three sites are collinear, no four sites are coplanar, etc. Given a vector in d dimensional space, we wish to find a $d-1$ dimensional hyperplane L normal to the vector such that an equal number of sites are on either side of the plane. (We use the following notation: space is divided by L into three disjoint connected regions L^+, L, and L^-, where L^+ denotes the open half-space on side of the plane L in the direction of the vector, and L^- is the half-space on the other side of L.) Clearly, if such an L exists, we can always find it in $O(n)$ time by projecting the sites onto a line parallel to the vector, finding a median of the point set (counting multiplicities) in $O(n)$ time, and drawing the corresponding plane.

Suppose we are given the sites in advance and are allowed to precompute information about them. Now, how well can we do? Given a query vector, how quickly can we find a median hyperplane normal to that vector? In this paper we address this problem and prove that the time required to search for the median of n sites is $O(\min(n, S(n^d, d-1)))$ with preprocessing time $O(n^d logn)$. Here $S(x, y)$ is the time required to perform deterministic geometric search; that is, given a y dimensional Euclidean region divided into x subregions by y dimensional simplices, the time required to find the subregion in which a given point lies. In fact, the geometric structure which our algorithm generates has a number of special properties which might facilitate a geometric search procedure that is more efficient than the most general case.

As specific results, for two and three dimensions we can find median separators in $O(\log n)$ time. This uses binary search for the two dimensional case, and any optimal two-dimensional search algorithm for the three dimensional case (such as persistent search trees [6]). This is an asymptotically optimal result.

A related problem is the *k-set* problem in which we wish to enumerate all separations of the set of sites into ordered pairs of a set with k sites and a set with $n-k$ sites. For each of these ordered pairs, there must exist a $d-1$ dimensional hyperplane L such that first half of the ordered pair lies entirely in L^+ and the second half lies entirely in L^-. Our techniques yields a new algorithm for the k-set problem for $d > 2$ and yields an alternative derivation of an algorithm due to Erdos, Lovasz, Simmons and Strauss for $d = 2$ [5]. Finding k-sets is known to be closely related to drawing k-order Voronoi diagrams [3].

Section 2 of the paper discusses the two dimensional problem and introduces some concepts useful for the d-dimensional case. The more general situation is addressed in Section 3.

2 Two Dimensions

2.1 Overview

We present an algorithm to solve the median location problem in two dimensions. The algorithm uses $O(n^2 \log n)$ time to preprocess n sites and can answer median queries in $O(\log n)$ time.

Let S be a set of n distinct sites in the plane in general position (i.e., no three collinear). We make the following definition:

Definition 2.1 (Median Line) *ℓ is a* median line *iff ℓ is a line and* $|\ell^+ \cap S| = |\ell^- \cap S|$.

Note that this definition allows sites of S to lie on ℓ. We are interested in the manner in which median lines partition S. Partitions of S are referred to as *separations.*

Definition 2.2 (Site-distinct) *Two median lines for S, ℓ_1 and ℓ_2, are said to be* site-distinct *iff*

$$\ell_1^+ \cap S \neq \ell_2^+ \cap S;$$

i.e., if they separate S differently.

Given an arbitrary angle $\phi \in [0, \pi)$, the task at hand is to specify a median line ℓ intersecting the x-axis at angle ϕ. (In general, we will say "line of angle ϕ" when referring to a line making angle ϕ with the x-axis.) We wish to answer such a *median line query* as fast as possible and we assume that many such queries will be asked for a particular choice of S. For the remainder of this section we assume that $|S|$ is even; the case where $|S|$ is odd is postponed until Section 2.4.

First we demonstrate that a solution to any given query must exist by presenting a naive algorithm to find one:

Lemma 2.3 *For every angle $\phi \in [0, \pi)$, there is a median line of angle ϕ.*

Proof: Pick an angle $\phi \in [0, \pi)$ and a line ℓ of angle $\phi + \frac{\pi}{2}$ (see Figure 1). Fix a coordinate axis (**x** or **y**) not parallel to ℓ and let F be the map taking points to the coordinate of their perpendicular projection onto ℓ. Observe that F^{-1} is a bijection between points on ℓ and lines of angle ϕ.

Project the sites in S onto ℓ and consider the multiset $F(S)$ of these projections. Find a median of $F(S)$—that is, a site a_1 such that there are exactly $\frac{n}{2} - 1$ sites $b \neq a_1$ with $F(b) \geq F(a_1)$. (Notice that this requires that n be even.) Find a site a_2 such that there are exactly $\frac{n}{2}$ sites $b \neq a_2$ with $F(b) \geq F(a_2)$.

Clearly $F(a_1) \geq F(a_2)$. If the inequality is strict, then for any point p on ℓ between $F(a_1)$ and $F(a_2)$ we know that $F^{-1}(p)$ is a median line of angle ϕ. Otherwise, let $p = F(a_1) = F(a_2)$. Then the line $F^{-1}(p)$ contains the sites a_1 and a_2 and can contain no others, since S is in general position. Hence it is a median line of angle ϕ. □

The algorithm described in Lemma 2.3 can be executed in $O(n)$ time by using a $O(n)$ median finding algorithm.

By examining the proof of Lemma 2.3 we can observe some other properties of median lines and separations. The first and most crucial is that all median lines of a given angle must separate S in the same manner (if they didn't, then what would happen to the sites between the lines?). As a result, we also know that if a median line contains a site it can be the only median line of that slope, and if a median

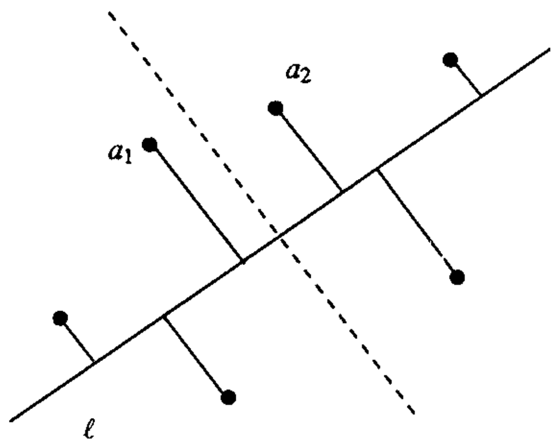

Figure 1: The dashed line is a median for the sites shown.

line contains no sites there can be infinitely many such median lines of that angle, none of which are site-distinct.

We show how preprocessing S may be used to reduce the 2-D median query problem to that of determining in which of a set of disjoint intervals a given query angle appears. The later problem can then be quickly solved using binary search. This allows us to answer median queries in $O(\log n)$ time per query, after preprocessing.

The algorithm for preprocessing S creates a binary tree containing triples of the form $\{\theta_i^{min}, \theta_i^{max}, P_i\}$, $1 \leq i \leq k$. This set of triples has the following properties:

(a) P_i is a set of exactly one or two points;

(b) each angle is in $[0, \pi]$ and each angle in $(0, \pi)$ is the angle of a line between two sites;

(c) $\theta_i^{min} < \theta_i^{max}$;

(d) $\theta_i^{max} = \theta_{i+1}^{min}$, $i \neq k$;

(e) $(\theta_j^{min}, \theta_j^{max}) \bigcap (\theta_i^{min}, \theta_i^{max}) = \emptyset$, $i \neq j$ (i.e., the intervals $(\theta_i^{min}, \theta_i^{max})$ are pairwise disjoint);

(f) $\bigcup_i [\theta_i^{min}, \theta_i^{max}] \subset [0, \pi]$;

(g) if $\phi \in [\theta_i^{min}, \theta_i^{max})$, then there is a constant time computation that yields a point $p \in P_i$ such that the line going through p having angle ϕ is a median line;

(h) if 0 is identified with π (and the last interval joined with the first if there is a horizontal median line going through any sites) and we consider the partition of $[0, \pi]$ into the endpoints and open intervals, then two medians are site-distinct if and only if their angles lie in different regions in this partition;

(i) there are at most n^2 such triples in the tree.

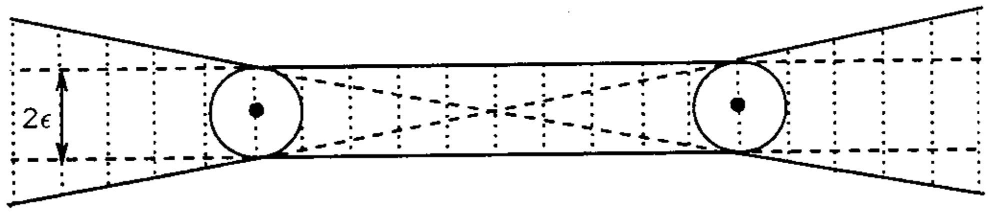

Figure 2: Circles have radius ϵ. No sites may be found between the solid lines shown, except for circle centers.

2.2 Preprocessing Step

We give the preprocessing step of our algorithm and show that it satisfies the requirements we have set out for it.

It will be helpful to think of replacing sites in S by balls of small radius centered at the site. We say that a line *abuts* a site if it is tangent to the associated ball. Similarly, talk about *rotating* a line about a site refers to rotating the line about this ball. In order for this change not to introduce any errors into the algorithm, we need for the radius ϵ of each ball to be small: if ℓ is a line that passes through two sites of S and we move ℓ so that it now only abuts those two sites, then no sites may have moved from ℓ^+ to ℓ or to ℓ^-, and similarly for ℓ^-; i.e., the separation of S has not changed (see Figure 2). Taking ϵ to be less than half the minimum distance from a site to the line connecting two other sites suffices. It takes $O(n^3)$ time to compute this value; however we stress that to discuss and implement the algorithm we only need to know that ϵ exists—an existence implied by the general position of S. We never need to do this $O(n^3)$ computation.

We also fix a positive number δ smaller than the distance between any two sites in S. δ is used to determine the P_i. The algorithm actually needs the value of δ, which can be computed in $O(n^2)$ time.

Our preprocessing algorithm generates, in order of increasing angle, a sequence of median lines that abut pairs of sites in S. We proceed as follows:

1. Use the method of Lemma 2.3 to find a horizontal median line ℓ. As the lemma shows, there are exactly two cases to consider:

   - ℓ contains two sites of S: call these sites s_1 and s_2 and perturb ℓ slightly to obtain a line ℓ_1, such that ℓ_1 abuts s_1 from the top and s_2 from the bottom (see Figure 3A).
   - ℓ contains no sites: let s_1 be the nearest site above ℓ and translate ℓ so that it abuts s_1 from below. Rotate ℓ counterclockwise about s_1 until it abuts on a second site s_2 and call this line ℓ_1 (see Figure 3B).

   In either case, let us call s_1 the *fixed site* of ℓ_1 and s_2 the *blocking site* of ℓ_1.

   We distinguish between two types of median lines.

   - The fixed and blocking sites lie on opposite sides of the line: we then say that the median is a *pivot line* (see Figure 3A).

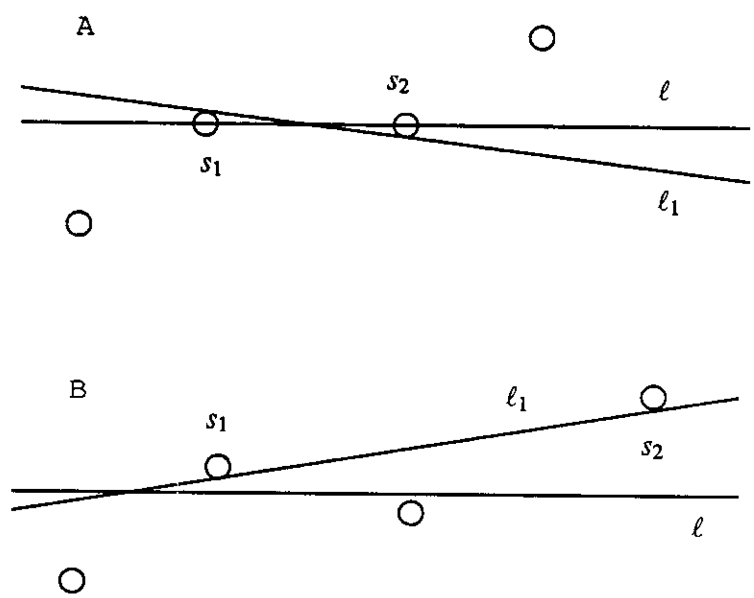

Figure 3: Possible positions of initial median lines.

- The fixed and blocking sites lie on the same side of the line: we then say that the median is a *swing line* (see Figure 3B).

The algorithm maintains the following *pivot invariant*:

- each ℓ_i is a median line that abuts two sites and if we rotate ℓ counterclockwise about the fixed site, the first site we bump into is the blocking site.

2. Suppose we have median line ℓ_i abutting two sites: fixed site s_1 and blocking site s_2. Obtain median ℓ_{i+1} with fixed site t_1 and blocking site t_2 from ℓ_i as follows:

   (a) If ℓ_i is a swing line, let ℓ_i' be ℓ_i and let $t_1 = s_2$ (see Figure 4A).

   (b) If ℓ_i is a pivot line, let ℓ_i' be the line obtained by rotating ℓ about the midpoint of $\overline{s_1 s_2}$ counterclockwise until ℓ is once again tangent to the ϵ-balls of the sites. By our choice of ϵ, the separation induced by ℓ_i' is the same as that induced by ℓ, except that s_1 and s_2 have switched sides. Let t_1 be either s_1 or s_2 (see Figure 4B).

   (c) Suppose ℓ_i' has angle θ. If there is no site t_2 such that $\overline{t_1 t_2}$ has angle in (θ, π), goto Step 2f. Otherwise, let t_2 be the site such that $\overline{t_1 t_2}$ has the smallest angle greater than θ. Copy ℓ_i' and rotate the copy counterclockwise about t_1 until it abuts t_2; let this line be ℓ_{i+1}.

   (d) If ℓ_i is the first pivot line that we have encountered so far, let θ_1^{min} be the angle of $\overline{s_1 s_2}$.

   (e) If ℓ_{i+1} is the mth pivot line found, $m > 1$, we must do two things:

      i. Let θ_{m-1}^{max} be the angle of $\overline{t_1 t_2}$, and let $\theta_m^{min} = \theta_{m-1}^{max}$.

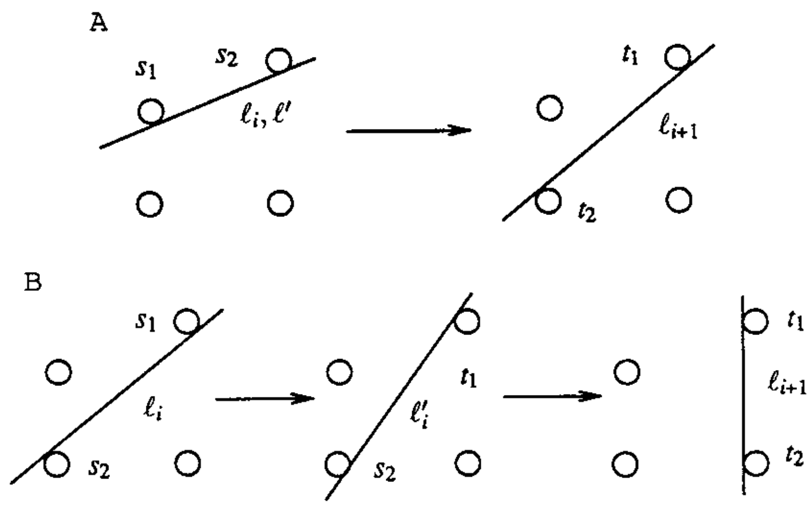

Figure 4: A: swing step; B: pivot step

ii. Let ℓ_j be the $m-$ 1st pivot line found. We compute P_{m-1} by considering two cases:
   - ℓ_j and ℓ_{i+1} intersect in a site s. Let b be the point on ℓ_j at distance δ from s and such that b is on the opposite side of ℓ_{i+1} from s. Pick point c in a similar manner on ℓ_{i+1}, opposite to s with respect to ℓ_j (see Figure 5). Let $P_{m-1} = \{b, c\}$.
   - ℓ_j and ℓ_{i+1} do not intersect in a site. Take P_{m-1} to be the point of intersection of the two lines.

(f) Now we have found all the pivot and swing lines. To account for the last part of $[0, \pi)$, let $\theta_k^{max} = \theta_1^{min} + \pi$, where k is the number of pivot lines found.

(g) Finally, put the triples into a balanced binary tree, ordered by the θ_i^{max}.

Lemma 2.4 *The ℓ_i are median lines.*

Proof: By construction, ℓ_1 is a median line and the pivot invariant holds. At every iteration of Step 2, if the pivot invariant holds, the next line generated must be a median line: in the swing case, the new line separates S in the same manner as the old; in the pivot case, the blocking site goes from one side to the other and the fixed site switches the other way, and all other sites remain unchanged. Since we generate lines in order of increasing slope, the pivot invariant holds after each execution of Step 2. □

Each of the ℓ_i considered by the algorithm is determined by two sites in S and each pair of sites determines at most one of the ℓ_i. There are only $O(n^2)$ pairs of sites, and so there can be at most this many ℓ_i. Since each step clearly terminates, our algorithm must halt and only $O(n^2)$ triples can have been generated.

Now that we know our algorithm halts, let's examine the triples generated. First, it is clear that $\theta_i^{min} < \theta_i^{max}$ and that the intervals generated are pairwise disjoint, with the possible exception of their endpoints. Also, any point of $[0, \pi)$ is in some interval, once we account for wraparound with respect to π.

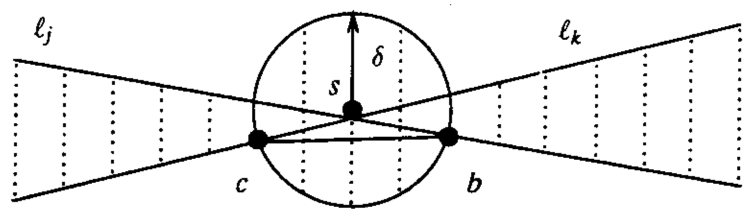

Figure 5: ℓ_i and ℓ_{i+1} abut site s. No points may lie in shaded area.

Lemma 2.5 *For each i, there is a median line of angle θ_i^{min} that passes through two sites.*

Proof: Fix i. θ_i^{min} corresponds to some pivot line ℓ. Let s and t be the blocking and fixed sites of ℓ. The line $\overline{st}$ is a median line: it separates S in the same manner as does ℓ, except for s and t, which were on different sides of ℓ, and are now on the line. By definition, $\overline{st}$ has angle θ_i^{min} and so is the required median line. □

Lemma 2.6 *If $\phi \in (\theta_i^{min}, \theta_i^{max})$, there is a median line of angle ϕ that passes through no sites. In fact, such a line goes through a point in P_i.*

Proof: Let ℓ_j be the pivot line corresponding (in the sense of the previous lemma) to θ_i^{min} and define ℓ_k similarly with respect to θ_i^{max}. Let $\ell_{j+1}, \ldots, \ell_{k-1}$, be the (possibly empty) sequence of swing lines that are generated by the algorithm between ℓ_j and ℓ_k.

Suppose that ℓ_j and ℓ_k do not abut a common site. Then P_i contains exactly one point—the intersection p of these lines. Each of the swing lines between ℓ_j and ℓ_k separate S in the same manner as ℓ_k since separations are only changed at pivot lines. This implies that the line ℓ_j' generated in Step 2a of the algorithm must also separate S in the same way as does ℓ_k. ℓ_j' and ℓ_k intersect in p (or arbitrarily close to p, as ϵ gets small). Imagine rotating ℓ_j' counterclockwise about p until it coincided with ℓ_k. If we crossed a site while doing this rotation, then ℓ_j' and ℓ_k would induce different separations of S, but we know that they separate S the same way! Hence, if $\phi \in (\theta_i^{min}, \theta_i^{max})$, then the line that has angle ϕ and passes through p (the point in P_i) is a median line and passes through no sites.

Alternatively, suppose ℓ_j and ℓ_k both abut on the same site s. Then the P_i we define will have two points. Define points b and c as we did in Step 2(e)ii of the algorithm (see Figure 5). Let θ be the angle of $\overline{bc}$. Note that $\theta_i^{min} < \theta < \theta_i^{max}$. The triangle determined by s, b and c lies inside the δ-ball about s and by definition of δ, no site is within δ of s. Hence, no site in S may lie in the interior of the triangle. By the same reasoning we used for the previous case, the region swept out by rotating a copy of ℓ_j' into ℓ_k must contain no sites. Let $\phi \in (\theta_i, \theta]$. Then the line of angle ϕ through b cannot intersect any site and must separate S in the same manner as ℓ_k and is therefore a median line that does not pass through a site. If $\phi \in [\theta, \theta_{i+1})$, the line through c at angle ϕ has the same properties. □

Corollary 2.7 *The θ_i^{min} are exactly the angles for which median lines pass through sites.*

The preceding lemmas characterize the angle intervals we have generated and established the properties we claimed of the triples. The definition of P_i gives us (a). Each θ^{min} and θ^{max} is the angle of a pivot line (except possibly for 0 and π, introduced by wraparound), hence (b). The sorting of triples by θ^{max} gives us (c) and (d); their construction—as the intervals between pivot lines—gives us (e) and (f). The constant time computation needed for (g) is outlined in the proof of Lemma 2.6. If two intersecting lines induce the same separation of S then the any line in the region swept out by rotating one into the other about the intersection induces the same separation, hence (h). The number of pairs of sites limits the number of pivot lines possible, hence (i).

Until now, we have been discussing this preprocessing step in terms of lines abutting ϵ-balls about the sites. When we actually implement the algorithm, we can dispense with this notion. To compute the θ_i^{min} and the θ_i^{max} we can simply work in terms of lines connecting sites keeping track of which site is the fixed site and which is the blocking site (conceptually). To compute the P_i we also need to know on which side of each pivot line (again, conceptually) the fixed and blocking site lie.

Now, let's see how long this preprocessing step takes:

- To find δ, look at all pairs of sites and let δ be one-half of the smallest distance between any such pair. This takes $O(n^2)$ time.
- Fix a site and calculate the slopes of all lines going through this site and each other one. Sort the slopes of these lines and store the information in a balanced binary tree named after the chosen site. Do this for each site. This takes $O(n \log n)$ time for each site, for a total of $O(n^2 \log n)$ time.
- Determine a horizontal median line. This takes $O(n)$ time using a linear time median algorithm and Lemma 2.3.
- Each iteration of Step 2 takes $O(\log n)$ time to find the next line by searching the slope data structure of the fixed site. Every time we find a pivot line we specify the P_i in constant time. We execute the loop at most $O(n^2)$ times, for a total time of $O(n^2 \log n)$.
- Finally, we put the $O(n^2)$ triples into a balanced binary tree. This takes a total of $O(n^2 \log n)$ time.

Adding all this up, we see that we need $O(n^2 \log n)$ time to do all the preprocessing.

In terms of space considerations the tree of triples we generate clearly uses $O(n^2)$ space. Using some earlier results [5, 4] we can bound the number of triples generated by $O(n^{3/2})$, and our space bound is therefore the same (this bound won't help our running time, since we also need $O(n^2 \log n)$ time to order the lines through each point).

2.3 Answering Queries

Given the data structure generated by the preprocessing step, and a median query for angle ϕ, the algorithm for answering this query is very straightforward:

1. search for the triple with $\theta_i^{min} \leq \phi < \theta_i^{max}$;
2. if the point set for this triple contains a single point, output that point;
3. if the point set for this triple contains a pair of points, compare ϕ to the angle of the line connecting the points and output the correct point, as specified in the proof of Lemma 2.6.

The algorithm is given an angle for input and requests a line as output; it provides a point instead. Knowing the angle of a line and a point it goes through is specifies the line uniquely; however, the algorithm can be directly modified to specify the line in another way—perhaps as an equation.

Using binary search, the lookup takes $O(\log n)$ time and the specification of the line takes constant time. Thus we can answer median queries in logarithmic time.

2.4 Two Dimensional Generalizations

This algorithm can be generalized in several ways while maintaining its two dimensional character. We consider several of these.

As described above, our algorithm applies only when $|S|$ is even. It is useful to consider what modifications must be made when $|S|$ is odd; these insights will be useful when we consider the d-dimensional case. The first and most important difference is that Lemma 2.3 would no longer hold: any median line for a odd-sized set of sites goes through a site. We can run the preprocessing step in the same way, except that P_i will now consist solely of the fixed site for each pivot line, and there are *no* median lines when ϕ is equal to one of the θ_i^{min}.

Other variations of the problem include relaxing the general position assumption and allowing sites to be collinear or even coincident. We can extend the algorithm to answer queries about *offset medians*—lines ℓ for which $|\ell^+ \cap S| - |\ell^- \cap S| = k$, for some nonzero k—using almost the same algorithm, except we run in $[0, 2\pi)$ rather than in $[0, \pi)$. The algorithm can also be adapted to generate k-sets—a set $T \subset S$ with $|T| = k$ and $\ell^+ \cap S = T$ for some ℓ.

Much work regarding k-sets is reported in the literature. Erdos, Lovasz, Simmons, Straus give a deterministic algorithm for the two-dimension case that is similar to our preprocessing step [5]. Clarkson gives a randomized algorithm for finding k-sets [Clark].

3 *d* Dimensions

3.1 Overview

What happens when we generalize the previous problem to d dimensions? We now have a set S of n sites in general position in d-space and we want to preprocess this information so that given the orientation of a d–1–dimensional hyperplane in d-space we can quickly identify a median hyperplane of that orientation.

Our algorithm is essentially a generalization the two-dimensional case and uses many results similar to those in Section 2. The preprocessing consists of partitioning the space of possible orientations of hyperplanes into nice regions: given an orientation and the region in this orientation space in which it lies, we can draw a median of that orientation in constant time. The query step then consists of searching this partition and actually producing the median. So as before, our algorithm preprocesses the set of sites and reduces the median query problem to a d–1–dimensional geometric search. Unfortunately for us, the cost of all known algorithms for geometric search will dominate the cost of the algorithm. This will be discussed further in Section 3.6.

First we need to determine what the space of hyperplane orientations is. (In this paper, "hyperplane" always refers to a d–1–dimensional subspace of d space; a "plane" always refers to a two dimensional subspace of d space.) The d-dimensional analog to the slope of a line is the unit vector normal to a hyperplane. The natural bijection between unit vectors and points on the surface of a unit d-sphere prompts us to make the following definition:

Definition 3.1 (Orientation Space) *We define* Orientation Space (OS) *to be the surface of the unit d-sphere.*

Sometimes we will abuse this definition and talk about points in OS as if they were vectors.

We want to be able to talk conveniently about hyperplanes, orientations of hyperplanes, and separations of sites induced by hyperplanes.

Definition 3.2 ($\mathcal{H}$, $\mathcal{M}$, $\mathcal{S}$) *Define the following maps:*

- $\mathcal{H}$: *hyperplanes* $\mapsto$ OS, *the many-one map taking each hyperplane to its orientation. (Mnemonic: $\mathcal{H}$yperplane orientation.)*
- $\mathcal{M}$: OS $\mapsto$ *sets of hyperplanes, the map from orientations to the set of median hyperplanes of that orientation. (Mnemonic: $\mathcal{M}$edians.)*
- $\mathcal{S}$: *hyperplanes* $\mapsto$ *separations, the map from hyperplanes to the separation of the sites they induce. (Mnemonic: $\mathcal{S}$eparation.)*

The astute reader might point out now that a hyperplane really has two normals—one for each side. This means that $\mathcal{H}$ is really multi-valued, its values being antipodal (that is, diametrically opposed) in OS. For the remainder of this section, we fix some direction as being positive and take $\mathcal{H}(P)$ to be the point lying in the positive half of OS. Since the partition of OS we generate will be symmetric about the origin, this distinction will not matter.

The maps $\mathcal{M}$ and $\mathcal{S}$ have some nice properties that carry over from two dimensions—among other things, $\mathcal{S} \circ \mathcal{M}$ is well-defined:

Lemma 3.3 *If a median hyperplane contains a site, it is the only median of that orientation. If a median contains no sites, there are infinitely many medians of that orientation. All medians of a given orientation yield the same separation of S.*

Proof: This intuition behind this proof is very similar to that behind Lemma 2.3. Suppose we're given an orientation α. Let P be any hyperplane of that orientation. The d–1–space determined by P is spanned by $d-1$ mutually perpendicular vectors $\mathbf{x_1}, \ldots, \mathbf{x_{d-1}}$ lying in P; these vectors, together with a vector $\mathbf{x_d}$ normal to P, are a basis for d-space. There is a 1–1 correspondence between the hyperplanes of orientation α and points on the $\mathbf{x_d}$ axis obtained by mapping each hyperplane to its intersection with the axis.

Slide P in the $\mathbf{x_d}$ direction until all sites of S lie on the positive side of P. Let Δ be the difference between the number of sites on the positive side and the number on the negative side of P. Δ initially has value n and can only decrease monotonically as P slides in the $\mathbf{x_d}$ direction. Consider the partition of the $\mathbf{x_d}$ axis into points and open intervals induced by projecting the sites of S perpendicularly onto it. Δ is fixed on each region, and on each takes a different value. Hence at most one interval can have $\Delta = 0$. The lemma follows. □

Our goal is to be able to compute quickly a map $\mathcal{M}'$ taking each point in OS to a representative median hyperplane (or to the empty set if no median exists). As in two dimensions, we proceed by partitioning OS into regions where $\mathcal{S} \circ \mathcal{M}$ is fixed. In two dimensions we did this by scanning all of OS and recognizing angles of pivot lines. This presented no difficulties since both a circle and time are one dimensional. For general d we are unable to search in this manner; we need to find the partition some other way.

An alternative method of constructing the partition for two dimensions is first to record all the angles corresponding to lines connecting pairs of sites and then to erase all the angles except the angles of pivot lines. (Recall that when $|S|$ is even, the pivots are exactly the medians containing sites.) This alternative procedure is generalized to solve the d dimensional problem. The next few sections develop the geometry of this partition, show that it has the appropriate qualities, and exhibit an $O(n^d \log n)$ algorithm for constructing it.

3.2 Geometry of Orientations of Site-Induced Hyperplanes

We first derive some general geometric results that apply to any problem involving hyperplane separators constrained to go through at least two sites.

To refer to such hyperplanes and their orientations we introduce the following:

Definition 3.4 (Basic Frame) *The* Basic Frame (BF) *is the subset of* OS *containing the orientations of all hyperplanes that contain at least two sites. Formally,*

$$\mathrm{BF} = \{\mathcal{H}(P) : P \text{ is a hyperplane } \wedge\ |P \cap S| \geq 2\}.$$

Definition 3.5 ($\mathcal{N}$) *Define* $\mathcal{N}$: BF $\mapsto$ *hyperplane;* $\mathcal{N}(\alpha) = P$ *where* P *is a hyperplane such that* $\mathcal{H}(P) = \alpha$ *and* $|P \cap S| \geq 2$. *(Mnemonic: the hyperplanes* $\mathcal{N}$*aturally associated with the* BF.*)*

We include a point α from *OS* is in the *BF* when there is hyperplane H_α of that orientation containing at least two sites of S. $\mathcal{N}$ takes each α in *BF* to its H_α. There is one flaw with this: $\mathcal{N}$ is not well-defined. There can exist two distinct subsets of S that are parallel; i.e., they are contained in hyperplanes with the same orientation. For the moment we disallow such S, and assert that $\mathcal{N}$ is single-valued; in Section 3.3 we show that this simplification is unnecessary.

The *BF* has a rich geometric structure that we now examine, but first a small note about geometric definitions. We will often talk about the number of dimensions of some subset of *OS* or the *BF*. This will always refer to the number of dimensions relative to the $d-1$-dimensional surface of the d-sphere, and not with respect to the d-dimensional space the hypersphere is embedded in. Thus, a one-dimensional subset of *OS* would be an arc.

Theorem 3.6 *Fix* k *distinct points in general position in* d*-space* $(2 \leq k \leq d)$. *Consider the collection of hyperplanes containing these points. The image in* OS *under* $\mathcal{H}$ *of that collection is a* $d-k+1$*-sphere, centered at the origin, and is thus a subsphere of* OS.

Proof: Let $\mathbf{n}$ be the unit normal to any hyperplane containing these points. Since these k points are in general position they determine a $k-1$-space. $\mathbf{n}$ is perpendicular to this space and so the normals of all these hyperplanes must lie in the $d-k+1$ dimensional space perpendicular to the $k-1$-space. Conversely, if $\mathbf{n}$ is a vector in this space, $\mathbf{n}$ must be the normal to some hyperplane containing the k points. Thus the set of orientations of the allowable hyperplanes is the intersection of a $d-k+1$ dimensional space and a d dimensional unit sphere (*OS*); i.e., a $d-k+1$ dimensional unit sphere. Since all of the vectors have unit length, this sphere is centered at the origin. □

This theorem shows that the *BF* consists of a collection of unit subspheres of *OS*, one for each $T \subset S$ where $2 \leq |T| \leq d$. This prompts us to define the following labeling method for spheres in *BF*:

Definition 3.7 (Subsphere Labeling) *If A is a subsphere in* BF *that corresponds to orientations of hyperplanes containing $T \subset S$, we call A the T*-subsphere.

Lemma 3.8 *If $U \subset T$, then the T-subsphere is contained in the U-subsphere.*

Proof: Any hyperplane containing the sites in T must contain the sites in U. □

Lemma 3.9 *Let $S_1,\ S_2 \subset S$ with $2 \leq |S_1|,\ |S_2| \leq d$. Then*

$$S_1\text{-subsphere} \cap S_2\text{-subsphere} = \begin{cases} \emptyset & \text{if } |S_1 \cup S_2| > d \\ S_1 \cup S_2\text{-subsphere} & \text{o.w.} \end{cases}$$

Proof: If $|S_1 \cup S_2| > d$ then no hyperplane can contain all the sites in both sets since the sites are in general position. Otherwise, the hyperplanes that contain all sites in both subsets are exactly the hyperplanes defining the $(S_1 \cup S_2)$-subsphere. □

We now want to partition the basic frame into its component pieces—points, edges, faces, etc. This presented no difficulty in two dimensions since it came for free: OS (the unit circle when $d = 2$) was divided by points (the angles of lines through sites).

Definition 3.10 (Σ) *Each point α in the basic frame is contained in at least one T-subsphere, for some $T \subset S$. Define $\Sigma(\alpha)$ to be the maximum of $|T|$, over all such T. (Mnemonic: the* count *of the elements in T).*

Lemma 3.11 *Each $\alpha \in$* BF *is contained in exactly one subsphere labeled with $\Sigma(\alpha)$ sites.*

Proof: Suppose $\alpha \in BF$ is contained in both the T_1 and T_2-subspheres, with $T_1 \neq T_2$ and $|T_1| = |T_2| = \Sigma(\alpha)$. Let $T = T_1 \cup T_2$. Then α contains all the sites in T, so the T-subsphere is non-empty. T properly contains the T_i, so we have $|T| > |T_i| = \Sigma(\alpha)$, contradicting the definition of Σ. □

We want to associate each $\alpha \in BF$ with the unique subsphere labeled with $\Sigma(\alpha)$ sites.

Definition 3.12 (Opened Subsphere) *Let the* opened T-subsphere *be the points α in the T-subsphere where $\Sigma(\alpha) = |T|$.*

Definition 3.13 (BF-pieces) *Let the* BF*-pieces be the maximal connected regions on the opened subspheres.*

These definitions are intended to capture the notion that higher dimensional spheres are divided into sections by lower dimensional spheres. Intuitively, the opened subspheres are obtained by removing all lower dimensional labeled subspheres from a given subsphere; i.e., pull out all points from the intersections of circles, pull out the circles from the intersections of spheres, and so on. If OS is the three dimensional sphere and we have some great circles and points of intersection of these circles, then the opened subspheres would consist of the points of intersection and each of the circles with these points removed. In this instance, the BF-pieces would consist of the points and the open arcs of the circles broken by the points.

Theorem 3.14 *The* BF*-pieces form a partition of the basic frame.*

Proof: Each point in the basic frame is contained in exactly one opened T-subsphere. $\square$

Theorem 3.15 *There are at most dn^d* BF*-pieces.*

Proof: Let a Σ-subsphere be an opened T-subsphere with $|T| = i$; let an i-piece be a maximal connected region on an i-subsphere. With this terminology, a point becomes a d-subsphere, and a circle is a $d-1$-subsphere.

The proof proceeds by induction.

Trivially, each d-subsphere is "broken" into one d-piece since a point can not be further divided.

Suppose each i-subsphere is broken into at most n^{d-i} i-pieces. Then each $i-1$-subsphere contains at most n^{d-i+1} i-pieces (since each $i-1$-subsphere contains at most n i-subspheres). Each of these i-pieces borders two $i-1$-pieces on the $i-1$-subsphere (one on the positive side and one on the negative), but each of these $i-1$-pieces is bordered by at least two i-pieces. Hence each $i-1$-subsphere is broken into at most $n^{d-(i-1)}$ $i-1$-pieces.

Inductively, for $2 \leq i \leq d$, each i-subsphere is broken into not more than n^{d-i} i-pieces. Since there are not more than n^i i-subspheres, this shows that there are at most n^d i-pieces in all. The result follows directly. $\square$

The reason for breaking the basic frame into these BF-pieces is that $\mathcal{N}$ provides a link between points in BF and hyperplanes: we will show that a transformation on the first corresponds to rotations in the other. In the development of this concept, the geometric notion of a *geodesic* on a hypersphere will be used. For our purposes, we may simply think of this as the intersection of some plane that passes through center of the hypersphere and the surface of the sphere. Alternatively, geodesics are great circles on the surface of the hypersphere. A *geodesic segment* is an arc on such a great circle.

Lemma 3.16 *Let α and β be distinct points on a geodesic segment* g *contained in BF. Then $\mathcal{N}(\alpha)$ and $\mathcal{N}(\beta)$ intersect in a d-2-dimensional subspace and the set of hyperplanes in $\mathcal{N}$(g) correspond to a rotation of $\mathcal{N}(\alpha)$ (about this subspace) into $\mathcal{N}(\beta)$.*

Proof: Since $\alpha \neq \beta$, $\mathcal{N}(\alpha)$ and $\mathcal{N}(\beta)$ cannot be parallel and must therefore intersect in a $d-2$-dimensional subspace. Let P be the plane defining **g**. The vectors α and β lie in P. This means that $\mathcal{N}(\alpha) \cap \mathcal{N}(\beta)$ is perpendicular to P. Pick any other point $\gamma \in$ **g**. Since $\gamma \in P$ the argument we used in Theorem 3.6 shows that $\mathcal{N}(\gamma) \cap \mathcal{N}(\alpha)$ is also a $d-2$-dimensional space perpendicular to P. However, P is two-dimensional and since we only have d dimensions in which to work, there can only be one $d-2$-dimensional space perpendicular to P. These two intersections must therefore be identical and changing an orientation from α to β only amounts to rotating $\mathcal{N}(\alpha)$ about this subspace. $\square$

Theorem 3.17 *$\mathcal{S} \circ \mathcal{N}$ is constant on each* BF*-piece.*

Proof: Suppose α and β lie in the same BF-piece on the opened T-subsphere. Since a BF-piece is a maximal connected region on an opened subsphere, there is a path **g** connecting them that also lies in the BF-piece. Let γ be any point on **g**. By the definition of BF-piece (Definition 3.4), $\mathcal{N}(\alpha)$, $\mathcal{N}(\beta)$ and $\mathcal{N}(\gamma)$ each contain the site t iff $t \in T$.

Suppose **g** is a geodesic path. Then the sites of T are in the subspace determined by the intersection of $\mathcal{N}(\alpha)$ and $\mathcal{N}(\beta)$. By Lemma 3.16 we can rotate the first of these hyperplanes into the second so that each hyperplane we get in the process contains this subspace, and hence contains T. Since each such

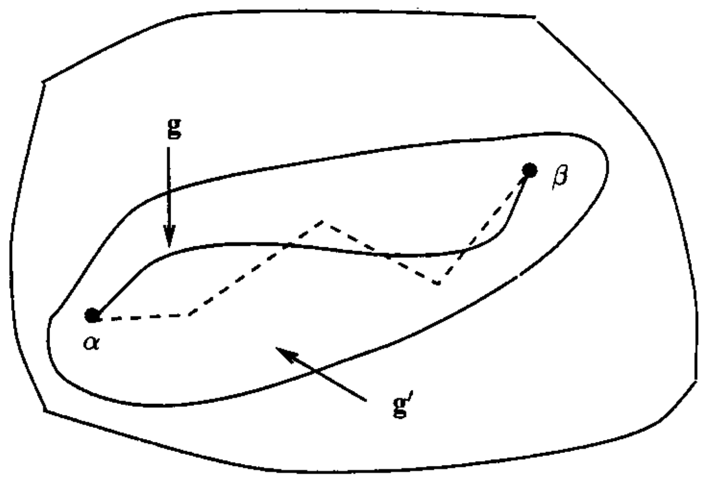

Figure 6: *BF*-piece is interior of region; **g** is non-geodesic; dashed line consists of geodesics inside **g**′.

hyperplane has orientation $\gamma \in \mathbf{g}$, these are the only sites that can be in the hyperplane. Therefore each of these hyperplanes induces the same separation of S, and so do $\mathcal{N}(\alpha)$ and $\mathcal{N}(\beta)$, and we are done.

Alternatively, suppose **g** is not a geodesic path. Then, since a one dimensional BF-piece is an arc of a great circle (i.e, a geodesic path), the *BF*-piece in which **g** lies cannot be one dimensional. A BF-piece is a connected open set. This means that we can expand **g** into an open region **g**′ such that $\mathbf{g} \subset \mathbf{g}'$ and **g**′ lies completely within the same BF-piece as **g** (for example, there is some ϵ such that we may take **g**′ to be the set of points in the BF-piece within ϵ of **g**). We may then connect α to β via a series of zig-zagging geodesics that all lie within **g**′—a piecewise geodesic path (see Figure 6). The earlier case shows that $\mathcal{S} \circ \mathcal{N}$ is constant on each closed geodesic segment of this path, and so this must hold for the whole path. □

3.3 Geometry of Orientations of Medians

The alternative two-dimensional algorithm described at the close of Section 3.1 first marked off the endpoints of regions in two-dimensional *OS* corresponding to lines containing two or more sites. It then deleted all of these regions except those that actually corresponded to medians (by Corollary 2.7 these were the pivot angles—exactly the slopes with unique medians).

In the previous section we performed the d-dimensional equivalent of the first part of this alternative algorithm. We now show how to perform the equivalent of the second part. First we develop "singularity" criteria for deciding which parts of the basic frame to throw out and which parts to keep and use for building our partition; we use a tool introduced in the proof of Lemma 3.3.

Recall that we have fixed a convention for the positive and negative sides of hyperplanes.

Definition 3.18 (Δ) *For a hyperplane P, define $\Delta(P)$ by*

$$\Delta(P) = |P^{+} \cap S| - |P^{-} \cap S|.$$

We shall write ΔP for $\Delta(P)$.

We can define Δ for points in the basic frame by considering $\Delta \circ \mathcal{N}$.

Lemma 3.19 *For $\alpha \in$ BF, $\mathcal{N}(\alpha) \in \mathcal{M}(\alpha)$ iff $\Delta\alpha = 0$.*

Proof: This follows directly from the definitions. □

In Section 2.4 we discussed generalizing the 2D algorithm to deal with the case when $|S|$ is odd. In this case Lemma 2.3 no longer holds—there *can* be angles for which no median exists. It turns out that in the higher dimensional case as well, but regardless of the parity of $|S|$, there can be orientations for which no median exists. We develop a series of results that show how to determine the *BF*-pieces such that there are no medians for orientations in that piece.

Lemma 3.20 *For $\alpha \in$ OS, $|\mathcal{M}(\alpha)| = 1$ iff $\alpha \in$ BF and $\Delta\alpha = 0$.*

Proof: This follows from the preceding lemma, and Lemma 3.3. □

Theorem 3.21 *For $\alpha \in$ BF, if $0 < |\Delta\alpha| < \Sigma(\alpha)$ then $\mathcal{M}(\alpha) = \emptyset$.*

Proof: We use the argument and notation used in the proof of Lemma 3.3. Assume α is as stated in the theorem. Since $\alpha \in BF$, the definition of Σ implies that we can slide P (a hyperplane of orientation α) to a point p on the $\mathbf{x_d}$-axis where P passes through $\Sigma(\alpha)$ sites. Δ takes on the value $\Delta\alpha$ at p. As we slide P off p in one direction, Δ will jump up by $\Sigma(\alpha)$; as we slide off in the other, Δ will jump down by the same amount. (Why? Because the $\Sigma(\alpha)$ sites that were on P all move simultaneously to one side of P.) Since Δ changes monotonically as we slide P along the $\mathbf{x_d}$-axis, on this axis Δ can only take on values in the range

$$(-\infty,\ \Delta\alpha - \Sigma(\alpha)] \cup \{\Delta\alpha\} \cup [\Delta\alpha + \Sigma(\alpha), \infty).$$

Since $0 < |\Delta\alpha| < \Sigma(\alpha)$, Δ can never take on the value zero, and there can be no median hyperplane for this orientation. □

Δ is fixed wherever Σ is, so the Δ value of a BF-piece is well-defined.

We now use this Δ criterion to pare down the BF.

Definition 3.22 (Refined Frame) *Let the* Refined Frame (RF) *be $\{\alpha \in \text{BF} : |\Delta\alpha| < \Sigma(\alpha)\}$; that is, the union (which is necessarily disjoint) of those BF-pieces whose Δ value has smaller magnitude than their Σ value.*

Recall that in the discussion following the definition of $\mathcal{N}$ (Definition 3.5) we observed that this map may in fact be multi-valued but that we would assume otherwise. This assumption was crucial in proving Lemma 3.11. We now dispense with this assumption. If we do not consider bad intersections of the subspheres comprising the *BF* (the ones that made $\mathcal{N}$ multivalued) to be valid links between *BF*-pieces, then we still can obtain the results from the previous section. The containment properties (Lemmas 3.8 and 3.9) of the subspheres still hold and the *BF*-pieces can be defined and have the same properties. So even if $\mathcal{N}$ is not well-defined, the *refined frame* is.

More, importantly, we have:

Theorem 3.23 *$\mathcal{N}$ is well-defined on the refined frame.*

Proof: To see this, recall the definitions used in the proof of Theorem 3.15. Suppose an i-piece *Alice* (of the T-subsphere, $i = |T|$) and a j-piece *Bob* (of the U-subsphere, $j = |U|$) have an unwanted intersection at α (that is, there is a hyperplane P containing T, and a hyperplane Q containing U, with $P \neq Q$ and $\mathcal{H}(P) = \mathcal{H}(Q) = \alpha \in Alice \cap Bob$). An axis projection argument similar to that used in the proof of Theorem 3.21 proves that $|\Delta Alice - \Delta Bob| \geq i+j$. Then, if $|\Delta Alice| < i$ we cannot also have $|\Delta Bob| < j$. Hence *Alice* and *Bob* cannot both be in the refined frame. □

We motivated the definition of the *RF* by claiming that it is a generalization of the two dimensional notion of pivot angles. In the following section we justify this claim by proving that the refined frame induces the desired partition of orientation space. In Section 3.5 we will then develop an algorithm to calculate the refined frame and the partition of *OS* it induces.

3.4 The Partition is Correct

Since the refined frame consists of *BF*-pieces with fewer than $d-1$ dimensions, it chops up the $d-1$-dimensional *OS* into chunks. We describe these chunks formally as the maximal connected regions in $OS \setminus RF$. These regions, together with the *BF*-pieces in the *RF*, constitute a partition of orientation space.

We prove that this partition is the one promised in Section 3.1.

Lemma 3.24 *If* $\alpha \in \mathrm{RF}$ *then* $|\mathcal{M}(\alpha)|$ *is zero or one.*

Proof: If $\Delta\alpha = 0$ then Lemmas 3.3 and 3.19 imply that $\mathcal{M}(\alpha) = \mathcal{N}(\alpha)$. If $\Delta\alpha \neq 0$ then the definition of *RF* allows us to apply Theorem 3.21, which tells us that $\mathcal{M}(\alpha) = \emptyset$. □

Lemma 3.25

$$\{\alpha : \alpha \in \mathrm{RF},\ \Delta\alpha \neq 0\} = \{\alpha : |\mathcal{M}(\alpha)| = 0\}.$$

If $|S|$ *is even, then also*

$$\{\alpha : \alpha \in \mathrm{RF},\ \Delta\alpha = 0\} = \{\alpha : |\mathcal{M}(\alpha)| = 1\}.$$

Proof: The first part follows from Theorem 3.21 and the definition of the *RF*. Recall that Lemma 3.3 implies that if a median passes through a site, it is the unique median of that orientation. If $|S|$ is even then, since all medians must pass through an even number of sites, medians passing through a site must pass through at least two. These medians thus have orientations that qualify for inclusion in the *RF*. The second part of the lemma follows. □

The following lemma gives us our main tool for solving our problem; it shows that we only need to consider two dimensional cases. In the next two results, we will make no distinction between the space in which S lies and that in which *OS* lies. We will define a plane intersecting *OS* in a particular way and then project S down to that plane. Since *OS* and S both lie in d-dimensional spaces, there is no problem doing this.

Lemma 3.26 (Geodesic Probing Lemma) *Consider an arbitrary great circle C in orientation space. Let C be the intersection of the plane P with the* OS *hypersphere. Let π be the function projecting points perpendicularly down to P. Then:*

- *C is the orientation space of lines lying in P;*

- *there is a one-to-one correspondence between $\mathcal{M}(C)$ and median lines for $\pi(S)$ lying in P.*

Proof: Let $\mathcal{H}_C$ be the set of all hyperplanes H such that $\mathcal{H}(H) \in C$.

For any hyperplane $H \in \mathcal{H}_C$, the normal to H lies in P and $H \cap P$ is a line perpendicular to this normal. This proves the first part of the lemma.

If $H \in \mathcal{H}_C$ is a median hyperplane for S, then $\pi(H)$ must be a median line for $\pi(S)$ in P, since π projects parallel to H. Alternatively, if ℓ is a median line for $\pi(S)$ in P, then $\pi^{-1}(\ell)$ is a median hyperplane whose orientation lies on C. This completes the proof. □

The main result of the previous section was Theorem 3.17 which showed that separations remained constant on BF-pieces. We extend that result to $OS\backslash RF$:

Theorem 3.27 *$\mathcal{S} \circ \mathcal{M}$ is constant on each connected region of* OS$\backslash$RF.

Proof: We will show that $\mathcal{S} \circ \mathcal{M}$ is constant on any connected geodesic segment contained in $OS\backslash RF$. Reasoning similar to that used in the proof of Theorem 3.17 then establishes the theorem.

Let R be some maximal connected region in $OS\backslash RF$. Let C be an arbitrary great circle that has nonempty intersection with R and let $C' = C \cap R$. C' will be some open arc (α, β), α, $\beta \in RF$.

Let P be the plane determined by C, as described by the Geodesic Probing Lemma; let π be the projection function onto P. Then (α, β) is also an interval of slopes in the orientation space of lines in P. Since median hyperplanes for S exist for every orientation in C' (because $C' \subset OS \setminus RF$), median lines for $\pi(S)$ exist for every slope in the slope interval C'. We want to show that these median lines all yield the same separation of $\pi(S)$; this will imply that the corresponding median hyperplanes all yield the same separation of S because we can just reverse the projection.

If the points of $\pi(S)$ are distinct and in general position in the two-dimensional space P, then this follows trivially from our work in two dimensions. The problem is that the points of $\pi(S)$ do not always meet these conditions. We have to broaden some of our two-dimensional work to handle such pathological site sets, where sites are not in general position and may even be coincident.

As we did in Section 2, consider small ϵ-balls around $\pi(S)$ (ϵ would be obtained the same way it would be obtained in the 2D case). Call these balls *siteballs*. Consider the problem of sweeping median lines around in this generalized two dimensional case. To obtain the needed result, we need to generalize our notion of swings and pivots. Swings work as before: if a line abuts several siteballs that all lie on the same side of the line, then we can swing the line in a positive direction around one of these siteballs and not cross any sites, and do the same in the negative direction about the other. Pivots still exist—a line abutting at least one siteball on each side—although what to do at a pivot step is no longer clear. This turns out to be irrelevant.

$\pi(\mathcal{N}(\alpha))$ and $\pi(\mathcal{N}(\beta))$ are two intersecting lines in P with orientations α and β; call these lines a and b and call their intersection c.

Let $\gamma \in C'$. Let ℓ be a median line in P of slope γ (we showed above that such an ℓ exists).

We consider the even and odd case for $|S|$ separately.

- If $|S|$ is even then ℓ contains no points in $\pi(S)$ (because otherwise γ would have to be on the refined frame). Translate ℓ until it abuts a siteball. Keep performing swing steps in the positive direction (counterclockwise in P) until you produce a pivot line; let A be the line *through* the sites this pivot abuts. Note that the pivot line (not A!) separates $\pi(S)$ in the same manner as ℓ since we crossed no sites during the swings. Similarly perform swing steps in the negative direction to obtain B. Let ΔA be the difference between the number of sites the pivot line abuts on one side and the number of sites it abuts on the other. Because this difference must be less than the number of sites abutted, $\mathcal{N}^{-1}(\pi^{-1}(A))$ must lie in the refined frame, and similarly for $\mathcal{N}^{-1}(\pi^{-1}(B))$. For orientations between that of A and B we have just swept ℓ through valid medians containing no sites, and hence $(\mathcal{N}^{-1}(\pi^{-1}(A)), \mathcal{N}^{-1}(\pi^{-1}(B)))$ must be the maximal connected segment of C containing γ.

Hence $a = A$ and $b = B$ (or $a = B$ and $b = A$). The same reasoning we used in Lemma 2.5 applies now: if there are no projected sites at c, then any line with orientation between those of a and b and drawn through c yields the same separation as ℓ, and hence is a median. If there are projected sites at c, we can draw medians of these orientations abutting the siteball there by making sure that all the sites are on the same side of the median.

- If $|S|$ is odd then ℓ must contain exactly one point in $\pi(S)$ (if it contained more, then γ would have to lie on the refined frame). We can obtain A and B by just rotating in either direction along this site until we hit another one. In doing so, we also sweep ℓ through valid medians containing exactly one site, so reasoning similar to that in the case above tells us that $(\mathcal{N}^{-1}(\pi^{-1}(A)), \mathcal{N}^{-1}(\pi^{-1}(B)))$ must be the maximal connected segment of C containing γ. So again, $a = A$ and $b = B$ (or the other way around); furthermore, any line in P of orientation between those of a and b through point c is a median for $\pi(S)$.

□

Theorem 3.28 *If two different maximal connected regions R_1 and R_2 in* OS\RF *are both bordered by the same* BF*-piece Alice with $\Sigma(Alice) = 2$, then Σ differs on R_1 and R_2.*

Proof: *Alice* is a common $d-2$ dimensional border of the $d-1$ dimensional regions R_1 and R_2. Draw a great circle C puncturing this border (i.e., C intersects each R_i in an open arc but intersects *Alice* in either one or two points). Apply the Geodesic Probing Lemma. □

The above work also suggests how we can compute $\mathcal{M}'$. Suppose we're given an orientation γ. If $\gamma \in RF$ then we're done: if $\Delta\gamma = 0$ then $\mathcal{N}(\gamma)$ is the only median and if $\Delta\gamma \neq 0$ there are none. Otherwise γ must lie in (the interior of) some region R (a maximal connected region in the partition of OS induced by the RF). Let C be a great circle (defined by plane P) passing through γ. C intersects the boundary of R on either side of γ at two points α and β. Let π again be the projection function taking things perpendicularly onto P. Let ℓ be the line tangent to the siteball at the intersection c of $\pi(\mathcal{N}(\alpha))$ and $\pi(\mathcal{N}(\beta))$. Then $\pi^{-1}(\ell)$ is our median hyperplane of orientation γ.

To make this approach work fully, we need to deal with this woolly business of ϵ-balls once again. We solved the analog of this problem in two dimensions by drawing a triangle around any site at the intersection of pivot lines. We can solve the generalized version of the problem by generalizing this solution. Basically, the idea is to construct a tetrahedral region near the intersection of median hyperplanes that abut a common site. If this region lies within the δ-ball (recall that δ is smaller than the minimum distance between sites in S) about c, this region will be guaranteed not to contain any sites.

More exactly, since S is in general position, there can be at most d sites that abut any hyperplane. We draw the d lines between these sites and s and label points on these lines that are at distance δ from s and on the other side of the hyperplane from s. Then, given an orientation in some region of OS, we can do $O(d)$ comparisons to figure out through which of these points to draw the median hyperplane.

3.5 The Algorithm

We will show how to construct a connection graph for the basic frame (indicating the details of the arrangement they form—which pieces border which other pieces, etc.) and compute the Δ values for the BF-pieces all in time $O(n^d \log n)$. This graph can be pared down to the refined frame and then expanded to include the maximal connected regions enclosed by the refined frame in time $O(n^d \log n)$.

We first need to examine the basic frame's structure a bit further. The BF is composed of BF-pieces and each BF-piece is associated with three fundamental parameters: its Δ value, the set $T \subset S$ indicating from which opened subsphere it came, and its Σ value ($|T|$). Each BF-piece with Σ-value $k < d$ can be thought of as an open $d-k$-dimensional region (part of some $d-k+1$-dimensional sphere) whose border consists of a collection of other BF-pieces with fewer dimensions (and hence with larger Σ values).

Definition 3.29 (Direct Border) *We will say that a* BF*-piece Alice* directly borders *a* BF*-piece Bob when*

- *Alice is part of the border of Bob;*
- $\Sigma(Alice) = \Sigma(Bob) + 1$*;*
- *if Alice is labeled with T and Bob is labeled with U, then* $U \subset T$ *and* $T \setminus U = \{s\}$*, where s is a single site.*

Definition 3.30 (∂, ∂^2) *Denote by* $\partial(Alice)$ *the set of* BF*-pieces that directly border the* BF*-piece Alice. Extend* ∂ *in the natural way to act on sets of* BF*-pieces. Denote by* $\partial^2(Alice)$ *the set* $\partial(\partial(Alice))$.

Each BF-piece with Σ-value k ($2 < k < d$) directly borders a set of BF-pieces with Σ-value $k-1$, and in turn is directly bordered by a set of BF-pieces with Σ-value $k+1$. If the BF-piece in question is labeled with set T, then the pieces it directly borders are each labeled with some set of the form $T \setminus \{s\}$ ($s \in T$), and the pieces it is directly bordered by are each labeled with some $T \cup \{s\}$ (for some $s \in S \setminus T$).

Definition 3.31 (Connection Graph) *Given a collection of* BF*-pieces, create a node for each one and draw directed edges from the node for Alice to the node for Bob when* $Bob \in \partial(Alice)$*. Label the nodes with their* BF*-pieces (including T-sets,* Σ *values, and* Δ *values). Call this graph the* connection graph *for that collection.*

Consider a BF-piece *Alice* when $d = 4$ and $\Sigma(Alice) = 2$. *Alice* is an open portion of a 3-sphere and is directly bordered by open arcs of great circles. These arcs do not intersect; if we expand each arc $Bob \in \partial(Alice)$ to include the pieces in $\partial(Bob)$, they do intersect. This inspires the following lemma:

Lemma 3.32 *For any* BF*-piece Alice with* $d - \Sigma(Alice) \geq 2$ *the set* $\partial(Alice) \cup \partial^2(Alice)$ *is connected.*

Proof: Let $k = d - \Sigma(Alice)$. The BF-pieces in $\partial(Alice)$ have dimension $k-1$ and those in $\partial^2(Alice)$ have dimension $k-2$; the rest of the pieces comprising the border of *Alice* have dimension less than $k-2$. The border of *Alice* is homeomorphic to the surface of a k-sphere, which is a $k-1$ dimensional manifold. Let f be such a homeomorphism. Then f will map the BF-pieces of dimension less than $k-2$ in the border to subspaces of dimension less than $k-2$. Removing a finite number of such subspaces still leaves the k-sphere's surface connected. □

This allows us to characterize *Alice* by the BF-pieces in $\partial(Alice)$. Furthermore, if we are examining a BF-piece *Bob* in $\partial(Alice)$ we can find another such BF-piece by examining the BF-pieces *Carl* such that $\Sigma(Carl) = \Sigma(Bob)$ and $\partial(Carl) \cap \partial(Bob)$ is nonempty. That is, there is another piece of $\partial(Alice)$ in $\partial^{-1}(\partial(Bob))$. In addition, we know that such pieces must be labeled with a superset of the set labeling *Alice*.

Suppose a *BF*-piece *Alice* has Σ-value k, is labeled with set T, and is directly bordered by a *BF*-piece *Bob* that is labeled with U. Suppose further that we have a connection graph for all the BF-pieces with Σ-values $k+1$ and $k+2$. Identify the nodes with the *BF*-pieces they represent. We can outline a naive traversal of this graph that marks a superset of $\partial(Alice)$. Let A be a set of nodes and initialize $A = \{Bob\}$. Add *Carl* to A if there is a node $Doug \in A$ such that for some node *Edna*, $Doug \to Edna$, $Carl \to Edna$, and the edge taking *Carl* to *Edna* is labeled with a superset of T. If we do this until A cannot be increased, then $\partial(Alice) \subset A$.

By considering Δ values, we can modify this traversal to compute $\partial(Alice)$ exactly. Given $T \subset S$ such that $2 \leq |T| < d$, let $s \in S \setminus T$ and let $U = T \cup \{s\}$. Then the U-subsphere has dimension one less than the T-subsphere and thus it partitions the T-subsphere into three regions: two open half-spheres and the U-subsphere itself.

Recall that we fixed a convention for the positive and negative sides of a hyperplane.

Lemma 3.33 *Let s be as above. Then the open half-spheres described above can be labeled R_1 and R_2 such that*

$$P \in \mathcal{N}(R_1) \;\Rightarrow\; s \in P^+$$

and

$$P \in \mathcal{N}(R_2) \;\Rightarrow\; s \in P^-.$$

Proof: This is because separations are constant on *BF*-pieces; Theorem 3.17. □

There are two sides of the U-subsphere relative to the T-subsphere; from this lemma we can label one of the sides "positive" and the other "negative." If $\mathbf{g}$ is a path on the T-subsphere crossing the opened U-subsphere at α, and $\Delta\alpha = k$, then just as $\mathbf{g}$ leaves α to the negative side Δ drops to $k-1$, and as $\mathbf{g}$ leaves to the positive side Δ goes to $k+1$.

This tells us the way Δ values on *BF*-pieces relate: if a *BF*-piece of the U-subsphere has Δ value k then it borders two *BF*-pieces on the T-subsphere: one with Δ value $k+1$ and one with Δ value $k-1$.

Returning to the last *Alice* and *Bob* example, suppose *Alice* (labeled with T) is directly bordered by *Bob* (labeled with $T \cup \{s\}$), and *Bob* is directly bordered by *Doug* (labeled with $T \cup \{s, t\}$) (see Figure 7). *Doug* represents an intersection of the $T \cup \{s\}$-subsphere with the $T \cup \{t\}$-subsphere. In our graph (see Figure 8) there are only three other edges coming into *Doug* labeled with a superset of T: two from *Carl* and *Edna*—labeled with $T \cup \{t\}$, and one from the other part of the $T \cup \{s\}$-subsphere. We know from geometry that only one of these *BF*-pieces is in $\partial(Alice)$: either *Carl* or *Edna*. How can we tell which one? Suppose *Alice* has Δ value m. Then *Bob*'s Δ value must differ from m by one—suppose it is $m-1$. Then *Alice* lies on the positive side of the $T \cup \{s\}$-subsphere (relative to the T-subsphere). Only one of *Carl* and *Edna* lies on that side—this is the one also in $\partial(Alice)$. *Carl* and *Edna* have Δ values one greater than and one less than *Doug*'s; the piece with the larger Δ value is the one we want. This piece borders two *BF*-pieces labeled with T; only one of these will have $\Delta = m$ and so that one must be *Alice*.

This suggests that we may want to label each edge $N_1 \to N_2$ in our connection graph with $\Delta N_1 - \Delta N_2$. Then a direct modification of the naive traversal outlined above would allow us to determine $\partial(Alice)$ exactly. Namely, let $A = \{Bob\}$. Add *Carl* to A if there is a node $Doug \in A$ such that for some node *Edna*:

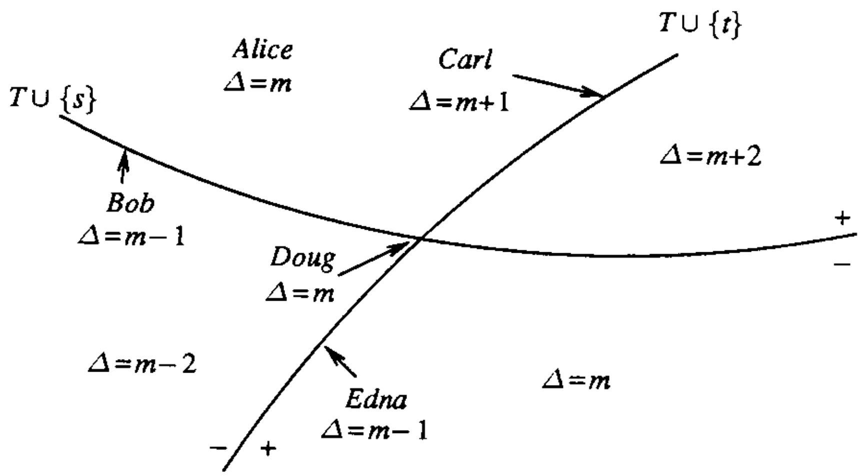

Figure 7: Surface of T-subsphere, $|T| = 2$, $d = 4$.

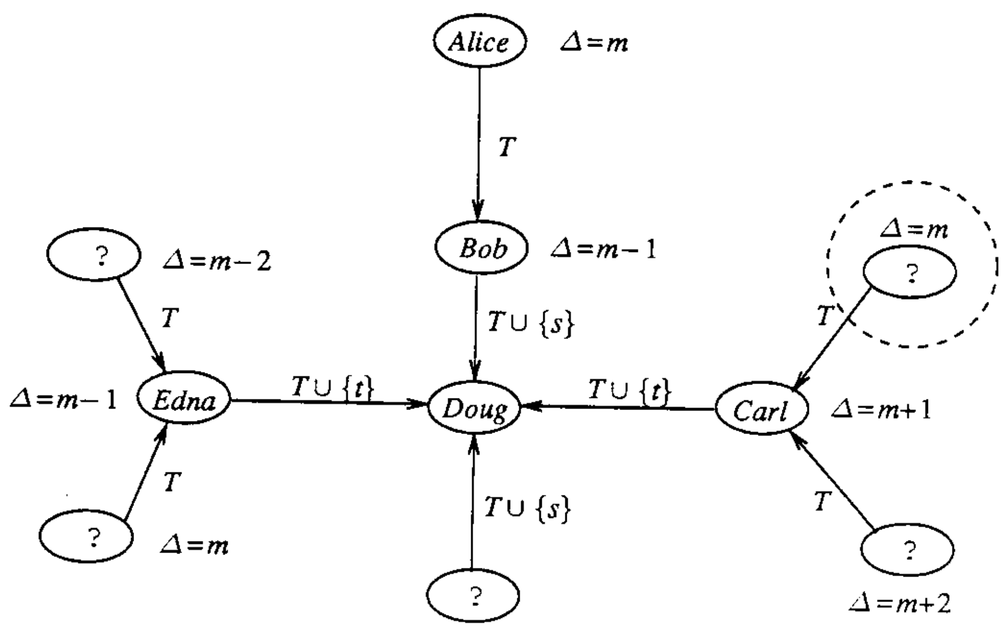

Figure 8: Connection graph for Figure 7. The circled node must be *Alice*.

1. *Doug* $\rightarrow$ *Edna* by edge e_1;
2. *Carl* $\rightarrow$ *Edna* by edge e_2;
3. e_1 and e_2 are labeled with different supersets of T;
4. the Δ difference on e_1 is the same as the Δ difference on e_2.

If we do this until A cannot be increased, then $\partial(Alice) = A$.

Our algorithm for constructing the connection graph of the basic frame proceeds as follows:

1. For each of the n^{d-1} subspheres whose Σ value is $d-1$ (these are great circles), sort (by angle in the plane of the circle) the occurrences of the contained n subspheres whose Σ value is d (these are points). (Total time: $O(n^d \log n)$.)
2. Each of the subspheres with Σ value $d-1$ is broken into BF-pieces by these contained d-subspheres. We can calculate Δ for one of those BF-pieces in time n. Since we have already sorted the occurrences of the d subspheres, we can then go around the circle and calculate Δ for the d-subspheres and the other BF-pieces in constant time per step. Construct nodes for each of the BF-pieces encountered so far, and draw and label the appropriate edges. There are $2n^d$ nodes, and $2n$ edges per each of the n^{d-1} great circles. Total steps: $O(n^d)$.

   We have now constructed the first two levels of our graph.
3. Let $i = d-1$. Consider each of the n^i BF-pieces with Σ value i. Let *Alice* be a typical one and suppose *Alice* is labeled with T and has Δ value k. Let T_s denote $T \setminus \{s\}$ for $s \in T$. There will ultimately be $2i$ edges coming into *Alice*: for each of the i T_s there is one with Δ value $k+1$ and one with Δ value $k-1$. Create these edges and the nodes at the other end (if they don't already exist).

   This may have created multiple nodes for *Alice*. We now remove the duplicate nodes and edges. Let *Carl* be the originating node of an edge going into *Alice*. For each node *Bob* such that *Alice* $\rightarrow$ *Bob*, go to *Bob* and then back up edges labeled with a superset of T_s other than T. Go up the T_s edge with the correct Δ value, and merge that node with *Carl* (i.e., all edges pointing to that node now point to *Carl* and all originating from that node now originate from *Carl*). Continue to do this with *Carl* until no more mergeable nodes are found. We can then do this with every other choice of *Carl* and then for every possible *Alice*. This entire operation is time bounded by the size of the graph, and so takes $O(n^d)$ time.
4. Repeat the above step for each $i \geq 2$ in descending order. This all takes time proportional to the number of edges, which is bounded by $2dn^d$.

To construct the connection graph for the refined frame, we just go through each node in the above graph and delete the ones whose Σ values exceed the magnitude of their Δ values. Again, this takes time proportional to the size of the graph, $O(n^d)$.

We also promised that we would include the d-dimensional maximal connected regions of the partition $OS \backslash RF$ in the connection graph. To determine these regions we only need to know how the BF-pieces having Σ-value 2 fit together. The problem is that these regions have no Δ values to guide us. Nevertheless, we can use one last bit of geometry to obtain the arrangement.

Denote by R the set of maximal connected regions in $OS\backslash RF$. Each BF-piece with Σ value 2 directly borders two regions in R—so for $r \in R$ we can define $\partial(r)$. Lemma 3.32 shows that if r is of dimension at least two then the union of all the BF-pieces in $\partial(r)$ and $\partial^2(r)$ is connected.

We will proceed as we did when building up the basic frame: enumerate $\partial(r)$ by moving from a member to one of its direct borders and then back to another member.

First, we need a geometric result.

Theorem 3.34 *Let Carl be a* BF*-piece with $\Sigma(Carl) = 3$. Let $Bob_1, ..., Bob_k$ be the* BF*-pieces that Carl directly borders (i.e., $\partial^{-1}(Carl) = \{Bob_i\}$). Denote by Bob_i^C the union of Bob_i with Carl. Then:*

1. $k \leq 12$.
2. *For any point $p \in Carl$ there is a unique plane P_p that intersects Carl exactly in p.*
3. *For any $p \in Carl$ and $i \leq k$, Bob_i^C intersects P_p in a half-open segment of a one-dimensional curve. The closed end of the segment has endpoint p. Denote this segment by $Bob_i^{C,p}$.*
4. *For any $p \in Carl$, the $Bob_i^{C,p}$ intersect precisely at p.*
5. *This allows us to order the $Bob_i^{C,p}$ by starting with $Bob_1^{C,p}$ and going clockwise (looking at P_p from some fixed orientation). This order is fixed, no matter what $p \in Carl$ is used.*

Proof:

1. *Carl* is labeled with a three-element set T. *Carl* can only directly border BF-pieces labeled with two-element subsets of T, and can only directly border two such BF-pieces for each of these subsets.
2. Fix p. At p, *Carl* is locally a $d-2$-manifold. Choose a basis for this local manifold—then there are two vectors $\mathbf{y_1}$ and $\mathbf{y_2}$ we can add to form a basis for d-space. Let P be the plane determined by the $\mathbf{y_i}$, originating at p.
3. At p, Bob_i^C is locally a $d-1$-manifold containing *Carl*. We can take the local basis for *Carl* found above and add some other vector $\mathbf{z}$ to form a local basis for Bob_i^C. $\mathbf{z}$ cannot be independent from $\{\mathbf{y_1}, \mathbf{y_2}\}$.
4. Any two Bob_i^C intersect exactly in C.
5. Suppose the orders at p_1 and p_2 differ. Let $\mathbf{g}$ be a path in *Carl* connecting these points. Then there must be a $p \in \mathbf{g}$ where two $Bob_i^{C,p}$ intersect in more than one point.

□

Let *Alice* be a BF-piece with $\Sigma(Alice) = 2$. She borders two regions r_1 and r_2 in R. The orientation we defined for hyperplanes can be extended to the $d-1$-dimensional face *Alice*. Each r_i lies on one side; let's say that r_1 lies on the positive side.

Given the connection graph for the refined frame, we can create nodes for r_1 and r_2 and draw edges coming into each such *Alice*. How do we merge the r_i representing the same regions? We simply calculate the order of the Bob_i for each BF-piece *Carl* with $\Sigma(Carl) = 3$. Then for each *Alice* we go to a bordering *Carl*, check how the positive/negative hyperplane orientations correspond to clockwise/counterclockwise in this order, and find new borders for *Alice*'s r_1 and r_2 one Bob_i away from *Alice* on either side. The orientation correlation also tells which region nodes of these new borders correspond to which of *Alice*'s

r_i. This process takes time proportional to the size of the connection graph for the refined frame, times a factor of $\log n$ lookup per step.

In all, we can then expand the connection graph for the refined frame to include the regions R in time $O(n^d \log n)$ and so the entire graph can be calculated in that time.

3.6 Search Considerations

We have a partition on a $d-1$ dimensional space (the OS surface) by points, arcs of great circles, parts of hyperspheres, and so on. To answer median queries, we need to identify in which region R a given point p in OS lies, and for an arbitrary great circle **g** through p, we need to find the segment of $\mathbf{g} \cap R$ in which p lies.

Standard geometric search techniques have not been designed for solving the point location problem on hyperspheres. Luckily, we do not such general techniques; normal Euclidean d-space algorithms suffice.

Theorem 3.35 *If the* BF *is mapped onto Euclidean $d-1$-space via a projection from the origin, then the resulting map consists entirely of points, lines, planes, hyperplanes, etc.; i.e., there are no curved pieces.*

Proof: We know from Theorem 3.6 that BF is composed of k-spheres, $1 \leq k \leq d-1$. Furthermore, these hyperspheres are centered at the origin. Fix some k-sphere S in BF. S is the intersection of a k-hyperplane passing through the origin and OS. The lines of projection from the origin also lie in this space, and conversely, every line in this space through the origin is the line of projection of some point in S. This means that the projected image of S in a $d-1$-space is the intersection of this space and the k-hyperplane just discussed. This is a $k-1$-dimensional space, proving the theorem. $\square$

Unfortunately, for $d > 3$ we know of no deterministic algorithm for geometric search in d–1 dimensions that takes $O(n^d \log n)$ preparation time and results in $O(\log n)$ query time. For $d = 3$, there are a number of recent algorithms [6] for doing this, so we have a complete algorithm for three dimensions.

Most research in point location deals with partitions more general then that produced by our preprocessing step. The partition we use has a great deal of geometric structure associated with it; perhaps this could be exploited make the search problem solvable within the time constraints we set. This is an area for further research.

4 Acknowledgments

We thank Victor Milenkovic for his assistance with the two-dimensional case.

J. Sipelstein was supported by a National Science Foundation Graduate Fellowship. S. W. Smith was supported by an Office of Naval Research Graduate Fellowship. J. D. Tygar was supported by a National Science Foundation Presidential Young Investigator grant.

The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the National Science Foundation or the US Government.

5 References

1. Chazelle, B., L. J. Guibas and D. T. Lee, "The Power of Geometric Duality." *Proc. 24th IEEE Annu. Symp. Found. Comput. Sci.* (1983) pp. 217–225.

2. Chazelle, B. and F. P. Preparata, "Halfspace Range Search: An Algorithmic Application of k-Sets." *Discrete Comput Geom* 1 (1986) pp. 83–93.

3. Clarkson, K.,"New Applications of Random Sampling in Computational Geometry." *Discrete Comput Geom* 2 (1987), pp. 195–222.

4. Edulsbrunner, H. and E. Welzl, "On the Number of Line Separations of a Finite Set in the Plane." *J. Combinatorial Theory* 38 (1985) pp. 15–29.

5. Erdos, P., L. Lovasz, A. Simmons, and E. G. Straus. "Dissection Graphs of Planar Point Sets," in *A Survey of Combinatorial Theory* (J. N. Srivastava et.al., Eds) North Holland, Amsterdam. 1973.

6. Sarnak, N. and R. E. Tarjan, "Planar Point Location Using Persistent Search Trees." *CACM* 29:7 (July 86), 669–679.

7. Welzl, E., "More on k-Sets of Finite Sets in the Plane." *Discrete Comput Geom* 1 (1986) pp. 95–100.