Centerpoint theorems for wedges

Jeff Erickson, Ferran Hurtado, Pat Morin

Discrete Mathematics and Theoretical Computer Science DMTCS vol. **11:1**, 2009, 45–54

Centerpoint Theorems for Wedges

Jeff Erickson[1] and Ferran Hurtado[2†] and Pat Morin[3‡]

[1] *Department of Computer Science, University of Illinois at Urbana-Champaign*
[2] *Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya*
[3] *School of Computer Science, Carleton University*

received July 3, 2008, *revised December 2, 2008*, *accepted November 11, 2008*.

The *Centerpoint Theorem* states that, for any set S of n points in $\mathbb{R}^d$, there exists a point p in $\mathbb{R}^d$ such that every closed halfspace containing p contains at least $\lceil n/(d+1)\rceil$ points of S. We consider generalizations of the Centerpoint Theorem in which halfspaces are replaced with wedges (cones) of angle α. In $\mathbb{R}^2$, we give bounds that are tight for all values of α and give an $O(n)$ time algorithm to find a point satisfying these bounds. We also give partial results for $\mathbb{R}^3$ and, more generally, $\mathbb{R}^d$.

Keywords: Centerpoint, Halfspace depth, wedges

1 Introduction

Let S be a set of n points in $\mathbb{R}^d$. The *halfspace depth* [14] of a point p with respect to S is defined as

$$D_\pi(p,S) = \min\{|h \cap S| : h \text{ is a closed halfspace that contains } p\} \enspace .$$

The *Centerpoint Theorem*, which is a simple consequence of Helly's Theorem [6], states that for any point set S of size n there exists a point whose halfspace depth is at least $\lceil n/(d+1)\rceil$. Furthermore, for every $n > 0$, there exists a point set S in $\mathbb{R}^d$ of size n for which no point in $\mathbb{R}^d$ has halfspace depth greater than $\lceil n/(d+1)\rceil$.

In this paper we consider a generalization of halfspace depth that we call α-*wedge depth*. Let r be a ray with endpoint q. An α-wedge with apex q and axis r is the point q plus the set of all points p such that the angle[(i)] between pq and r is at most $\alpha/2$. The α-wedge depth of a point p with respect to a point set S is defined as

$$D_\alpha(p,S) = \min\{|h \cap S| : h \text{ is an } \alpha\text{-wedge with apex } p\} \enspace .$$

†Partially supported by projects MECMTM2006-01267 and Gen. Cat. 2005SGR00692.
‡Partially supported by grants from NSERC, the Ontario Ministry of Research and Innovation, and the E. Bower Carty Memorial Foundation.

(i) We use the convention that the angle between two line segments (or, in this case, a ray and a line segment) with an endpoint in common is the smaller of the two angles occuring at the common point.

Several authors have studied α-wedge depth and related notions. The set of point in S with α-wedge depth 1 are called *unoriented α-maxima* by Avis *et al* [2] who study their computational and combinatorial properties. Abellenas *et al* [1] study (α, k)-sets in the plane. These are the subsets of S that can be separated from the remainder of S by an α-wedge. In particular, every (α, k) set defines a locus of points whose α-wedge depth is at most k. Several authors have studied the use of α-floodlights (α-wedges) for illuminating regions of the plane [7, 13].

In the current paper, we consider bounds on the points of maximum α-wedge depth. Define the function $f_\alpha^d(n)$ as follows:

$$f_\alpha^d(n) = \min\left\{\max\{D_\alpha(p,S) : p \in \mathbb{R}^d\} : S \subseteq \mathbb{R}^d\,,\, |S| = n\right\}$$

That is, f_α^d defines, for each n, the maximum value k for which every point set S of size n is guaranteed to define a point whose α-wedge depth with respect to S is at least k. The Centerpoint Theorem states that $f_\pi^d(n) = \lceil n/(d+1) \rceil$. In this paper we prove the following Theorem about 2-dimensional point sets:

Theorem 1

$$f_\alpha^2(n) = \begin{cases} 1 & \text{if } \alpha < \pi \\ \lceil n/3 \rceil & \text{if } \pi \leqslant \alpha < 4\pi/3 \\ \lceil n/2 \rceil & \text{if } 4\pi/3 \leqslant \alpha < 2\pi \\ n & \text{if } \alpha = 2\pi \;. \end{cases}$$

Furthermore, for any α and any point set S of size n, a point p such that $D_\alpha(p,S) \geqslant f_\alpha^2(n)$ can be found in $O(n)$ expected time.

We also prove some partial results about f_α^d for dimensions $d \geqslant 3$. The remainder of the paper is organized as follows. In Section 2 we fully characterize f_α^2. In Section 3 we give a partial characterization of f_α^d. In Section 4 we refine this characterization for the special case $d = 3$. Finally, in Section 5 we summarize and conclude with open problems.

2 Proof of Theorem 1

In this section we prove a sequence of lemmata that immediately imply Theorem 1.

Lemma 1 *If $\alpha < \pi$ then $f_\alpha^2(n) = 1$ and a point p such that $D_\alpha(p,S) \geqslant 1$ can be found in $O(1)$ time.*

Proof: To prove the lower bound, we observe that for any non-empty point set S, every point $p \in S$ satisfies $D_\alpha(p,S) \geqslant 1$, so $f_\alpha^2(n) \geqslant 1$. This proves the lower bound and gives an $O(1)$ time algorithm for finding p.

For the upper bound, consider a set S of points that are all on the x-axis. For any point p on or above the x axis, the α-wedge whose axis is vertical and upwards intersects the x axis in at most one point, therefore $D_\alpha(p,S) \leqslant 1$. For any point p below the x axis, the α-wedge whose axis is vertical and downwards does not intersect the x axis at all, so $D_\alpha(p,S) = 0$. In either case, $D_\alpha(p,S) \leqslant 1$ so $f_\alpha^2(n) \leqslant 1$. $\square$

Lemma 2 *If $\pi \leqslant \alpha < 4\pi/3$ then $f_\alpha^2(n) = \lceil n/3 \rceil$ and a point p such that $D_\alpha(p,S) \geqslant \lceil n/3 \rceil$ can be found in $O(n)$ time.*

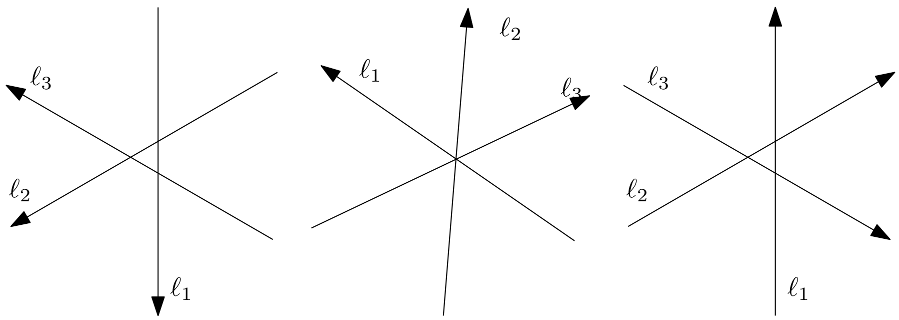

(a) (b) (c)

Fig. 1: The existence of three concurrent halving lines that meet at angles of $\pi/3$.

Proof: For the lower bound, we observe that every α-wedge containing p also contains a halfspace containing p. Therefore, the Centerpoint Theorem implies that $f^2_\alpha(n) \geqslant \lceil n/3 \rceil$. This proves the lower bound and the algorithm of Jadhav and Mukhopadhyay [10] gives an $O(n)$ time algorithm for finding p.

For the upper bound, consider the following point set. Start with three rays originating at the origin such that each pair of rays meet at an angle of $2\pi/3$. Place $\lceil n/3 \rceil$ or $\lfloor n/3 \rfloor$ points on each ray, as appropriate, so that the total number of points is n. For any point $p \in \mathbb{R}^2$, there exists a $4\pi/3$-wedge whose apex is at p and whose interior intersects only one of the three rays (the axis of this wedge is parallel to this ray). This $4\pi/3$ wedge contains an α-wedge that contains p and intersects only one of the three rays, therefore $D_\alpha(p, S) \leqslant \lceil n/3 \rceil$. Since the choice of p is arbitrary, this implies that $f^2_\alpha(n) \leqslant \lceil n/3 \rceil$. □

The next part of the proof uses the notion of halving lines. A *halving line in direction* d of a finite point set S, $|S| = n$, is a line ℓ parallel to d such that each of the 2 closed halfplanes bounded by ℓ contains at least $\lceil n/2 \rceil$ points of S. We will use the convention that, if n is even, then the closest point of S to the left of ℓ is at the same distance from ℓ as the closest point of S to the right of ℓ. In this way, a halving line is uniquely defined by its direction. The following lemma was proven by Fekete and Meijer [8, Lemma 2] in a different context. However, for completeness, we include a proof because an understanding of the existence proof is required for the algorithm described in Lemma 4.

Lemma 3 *For any point set S there exists three concurrent halving lines of S such that the angle*(ii) *between any pair of lines is $\pi/3$.*

Proof: To prove the existence of these three halving lines we start with one vertical halving line, ℓ_1, and the other two halving lines, ℓ_2 and ℓ_3, forming angles of $\pi/3$ with ℓ_1, ℓ_2 having positive slope and ℓ_3 having negative slope (Figure 1.a). If these three halving lines are concurrent then the construction is complete.

Otherwise, assume without loss of generality that ℓ_1 is directed downwards and that $\ell_2 \cap \ell_3$ is to its right. Imagine continuously rotating the three lines while maintaining the invariant that they are all halving lines and that the angle between any two is $\pi/3$. After having rotated the lines by an angle of π, the three halving lines are identical to their initial configuration except that the direction of ℓ_1 is reversed, so now

(ii) We use the convention that the angle between a pair of lines is the smaller of the two angles defined by the two lines.

$\ell_2 \cap \ell_3$ is to the left of ℓ_1 (Figure 1.c). We conclude that at some point during this process $\ell_2 \cap \ell_3$ must have been on ℓ_1 (Figure 1.b), at which point the three lines were concurrent. This completes proof. $\square$

Lemma 4 *Three halving lines satisfying the conditions of Lemma 3 can be found in $O(n)$ time.*

Proof: To find the three halving lines we apply the prune-and-search paradigm in much the same way as the algorithm of Lo, Matoušek, and Steiger [11] for finding planar ham-sandwich cuts. By the standard "computational geometry duality" [5, Section 1.3.3], our problem is to find three points on the median level of n lines such that these points are collinear and their x-coordinates satisfy a certain equation.

More precisely, given a set S^* of n lines (that are dual to the points of S), let

$$h_k(x) = \min\{y : (x,y) \text{ is on or above at least } k \text{ lines of } S^*\}$$

and let $h = h_{\lceil n/2 \rceil}$. The set of all points (x, y) satisfying $y = h_k(x)$ is called the *k-level* of S^* or, for $k = \lceil n/2 \rceil$, the *median level*. The dual of our problem is to find a value x such that the three points $(x, h(x))$, $(g_1(x), h(g_1(x)))$ and $(g_2(x), h(g_2(x)))$ are collinear. Here $g_1(x) = \tan(\arctan(x) + \pi/3)$ and $g_2(x) = \tan(\arctan(x) - \pi/3)$ which captures the condition that each pair of halving lines form an angle of $\pi/3$. [Informally, the continuity argument in the proof of Lemma 3 is equivalent to the observation that, if the sequence of points $\langle(-\infty, h(-\infty)), (g_1(-\infty), h(g_1(-\infty)), (g_2(-\infty), h(g_2(-\infty))\rangle$ form a right (respectively left) turn then the points $\langle(\infty, h(\infty)), (g_1(\infty), h(g_1(\infty)), (g_2(\infty), h(g_2(\infty))\rangle$ form a left (respectively right) turn, so there must be some $x \in (-\infty, \infty)$ such that $(x, h(x))$, $(g_1(x), h(g_1(x)))$ and $(g_2(x), h(g_2(x)))$ are collinear.]

Each iteration in the algorithm of Lo *et al* [11] constructs, in time linear in $|S^*|$, a trapezoid T that is guaranteed to contain a ham-sandwich point(iii) and that intersects at most $2n/3$ lines of S^*. The lines in S^* not intersecting T are then discarded and the algorithm recurses on the remaining lines. Since a constant fraction of the lines are discarded in each iteration, the running time of the algorithm is a geometrically decreasing series and is therefore $O(|S^*|)$.

In our setting, we are searching for 3 points, so at each iteration we construct three trapezoids T, T_1 and T_2 such that each trapezoid intersects at most δm lines, for an arbitrarily small constant $\delta < 1/3$. We then discard from S^* any line not intersecting any of the three trapezoids and recurse on the remaining lines. Each iteration (described below) takes $O(|S^*|)$ time and decreases the size of S^* by a factor of 3δ, so the entire algorithm runs in $O(|S^*|) = O(n)$ time.

Because the algorithm is recursive the subproblems it solves are slightly more general than the original problem. Given a set S^* of lines, two x-coordinates x_1 and x_2 and three integers k, k_1 and k_2, the algorithm finds an x-coordinate $x \in [x_1, x_2]$ such that the three points $(x, h_k(x))$, $(g_1(x), h_{k_1}(g_1(x)))$ and $(g_2(x), h_{k_2}(g_2(x)))$ are collinear. Such a value x is guaranteed *a priori* to exist. Note that, for our initial recursive call we set $x_1 = -\infty$, $x_2 = \infty$, and $k = k_1 = k_2 = \lceil n/2 \rceil$.

All that remains is to show how to implement a single iteration of the algorithm in $O(|S^*|)$ time. To begin, we create a set X of x-coordinates that initially contains the values x_1 and x_2. Next we add to X an additional $O(1)$ values so that, for any two consecutive elements of X, the arrangement of our m lines contains at most $(\delta m)^2/16$ vertices that have x-coordinates between these two elements of X. These additional values can be found in $O(|S^*|)$ time using (e.g.) the algorithm of Matoušek [12] (or much more simply by random sampling). Finally, for each value $x \in X$ we add the values $g_1^{-1}(x)$ and $g_2^{-1}(x)$

(iii) A ham-sandwich point is the dual of a ham-sandwich line.

to X. This last step guarantees that, for any two consecutive elements x'_1 and x'_2 of X, the arrangement of the lines in S^* contains at most $(\delta m)^2/16$ vertices whose x coordinates are in the range $[g_1(x'_1), g_1(x'_2)]$ (respectively $[g_2(x'_1), g_2(x'_2)]$).

Now, using $O(|X|) = O(1)$ applications of a linear time selection algorithm (e.g., [3]) we can find, in $O(|S^*|)$ time, two consecutive elements x'_1 and x'_2 of X such that $x'_1, x'_2 \in [x_1, x_2]$ and a solution to our problem lies in the interval $[x'_1, x'_2]$. Consider the trapezoid T whose four corners are given by $(x'_1, h_{k\pm\lfloor\delta m/4\rfloor}(x'_1))$ and $(x'_2, h_{k\pm\lfloor\delta m/4\rfloor}(x'_2))$. A simple argument [11] shows that this trapezoid intersects at most δm lines of S^* and that the k-level of S^* does not intersect the top or bottom sides of this trapezoid. Similarly, there are trapezoids T_1 and T_2 defined by the four points $(g_1(x'_1), h_{k_1\pm\lfloor\delta m/4\rfloor}(g_1(x'_1)))$ and $(g_1(x'_2), h_{k_1\pm\lfloor\delta m/4\rfloor}(g_1(x'_2)))$ and the four points $(g_2(x'_1), h_{k_1\pm\lfloor\delta m/4\rfloor}(g_2(x'_1)))$ and $(g_2(x'_2), h_{k_1\pm\lfloor\delta m/4\rfloor}(g_2(x'_2)))$, respectively. The inclusion of the elements of the form $g_1^{-1}(x)$ and $g_2^{-1}(x)$ in the set X guarantees that, neither T_1 nor T_2 intersect more than δm lines in S^* and the k_1-level (respectively k_2-level) of S^* does not intersect the top or bottom sides of T_1 (respectively T_2).

Altogether, this means that there are at least $m - 3\delta m$ lines in S^* that do not intersect any of the trapezoids T, T_1 or T_2. When we recurse, we discard these lines, set $x_1 = x'_1$, $x_2 = x'_2$, and subtract from k (respectively k_1 and k_2) the number of discarded lines that pass below T (respectively T_1 and T_2). This completes the description of the algorithm and the proof of the lemma. □

Lemma 5 *If $4\pi/3 \leqslant \alpha < 2\pi$ then $f^2_\alpha(n) = \lceil n/2\rceil$ and a point p such that $D_\alpha(p, S) \geqslant \lceil n/2\rceil$ can be found in $O(n)$ time.*

Proof: For the lower bound, consider the three halving lines whose existence is given by Lemma 3. These three halving lines naturally define six $\pi/3$-wedges. Observe that if we take p to be the common intersection point of the three halving lines then any α-wedge with apex p contains at least 3 consecutive $\pi/3$-wedges and therefore contains at least $\lceil n/2\rceil$ points of S Therefore, $f^2_\alpha(n) \geqslant \lceil n/2\rceil$, and the point p such that $D_\alpha(p, S) \geqslant \lceil n/2\rceil$ can be found in $O(n)$ time using Lemma 4.

For the upper bound, we consider a point set in which the points have been clustered into two groups of size $\lfloor n/2\rfloor$ and $\lceil n/2\rceil$. Each of the two groups is contained in a unit ball and the distance between the two groups is very large, say r. Now, observe that any point $p \in \mathbb{R}^2$ must be at distance at least $r/2$ from at least one of the two groups. This means that, if r is sufficiently large, then there exists a $(2\pi - \alpha)$-wedge whose apex is p and that contains this group in its interior. The complementary α-wedge contains p and does not contain any points of this group. Therefore, $D_\alpha(p, S) \leqslant \lceil n/2\rceil$. Since p was chosen arbitrarily, we conclude that $f^2_\alpha(n) \leqslant \lceil n/2\rceil$. □

Proof of Theorem 1: The theorem follows immediately from Lemma 1, Lemma 2 and Lemma 5. □

3 Some Results for $\mathbb{R}^d$

In this section, we consider α-wedge depth in $\mathbb{R}^d$, and prove some bounds on the function f^d_α. The following lemma results from exactly the same arguments used in the proofs of Lemma 1, Lemma 2 and Lemma 5 (namely points on a line, the Centerpoint Theorem, and 2 small clusters of points, respectively).

Lemma 6 *f_d^α satisfies the following:*

$$\begin{array}{lcll} f_\alpha^d(n) & = & 1 & \text{if } \alpha < \pi \\ f_\alpha^d(n) & \geqslant & \lceil n/(d+1) \rceil & \text{if } \alpha \geqslant \pi \\ f_\alpha^d(n) & \leqslant & \lceil n/2 \rceil & \text{if } \alpha < 2\pi \end{array}$$

The following technical lemma is needed for proving an upper bound that generalizes the construction in Lemma 2.

Lemma 7 *Let T be a regular d-simplex whose center is at the origin. Then, for any d vertices of T, there is a $2\arccos(1/d)$-wedge whose apex is at the origin and that contains these d vertices of S.*

Proof: Without loss of generality, we can consider the regular d-simplex whose vertices are $e_1, \dots, e_d, ((1-\sqrt{d+1})/d)(e_1 + \cdots + e_d)$ where e_i is the ith coordinate unit vector in $\mathbb{R}^d$. The center of this simplex is the point $c = \sqrt{d+1}/(d^2+d)(e_1 + \cdots + e_d)$. Consider the ray r that originates at c and contains the point $e_1 + \cdots + e_d$. The angle between r and e_i, for any $1 \leqslant i \leqslant d$ is easily determined to be $\arccos(1/d)$ using the famous formula

$$\measuredangle uv = \arccos\left(\frac{u \cdot v}{\|u\|\|v\|}\right)$$

for the angle between two vectors u and v. Thus, the d-vertices $e_1, \dots e_d$ are contained in the $2\arccos(1/d)$-wedge whose axis is r, as required. $\square$

The next lemma is a generalization of Lemma 2. Notice that $2(\pi - \arccos(1/d))$ approaches π from above as d increases. This means that, for sufficiently large d, the upper bound in the following lemma only holds for $\alpha < \pi + \epsilon$.

Lemma 8 *If $\alpha < 2(\pi - \arccos(1/d))$ then $f_\alpha^d(n) \leqslant \lceil n/(d+1) \rceil$.*

Proof: We use a generalization of the point set used in the proof of Lemma 2. Let T be a regular d-simplex whose center is at the origin and consider the $d+1$ rays originating at the origin and each containing a different vertex of T. On each of these rays, place $\lceil n/(d+1) \rceil$ or $\lfloor n/(d+1) \rfloor$ points, as appropriate, to produce a point set S of size n. We claim, as in the proof of Lemma 2, that for any point $p \in \mathbb{R}^2$, there is a $2\arccos(1/d)$-wedge whose apex is p and that contains d of the $d+1$ rays that contain the points of S.

To see why this is so, let $C_1, \dots, C_{d+1}$ be the closed cones obtained by taking the conical hull(iv) of each facet of T. Notice that these cones cover $\mathbb{R}^d$ and that each cone contains d of the $d+1$ rays that contain S. Furthermore, if the cone C_i contains the point $-p$ then, by Lemma 7, there is a $2\arccos(1/d)$-wedge whose apex is at p and that contains C_i.

If we consider the complementary $2(\pi - \arccos(1/d))$-wedge then the interior of this wedge does not intersect C_i and hence intersects only 1 of the $d+1$ rays that contain S. This $2(\pi - \arccos(1/d))$-wedge contains an α-wedge that contains p and contains at most $\lceil n/(d+1) \rceil$ points of S, as required. $\square$

Next we consider lower bounds. The following lemma, which is a generalization of a 3-dimensional result of Fekete and Meijer [8] is used to find centerpoints.

Lemma 9 *If $\alpha \geqslant \pi + 2\arccos(1/\sqrt{d})$ then $f_\alpha^d(n) \geqslant \lceil n/2 \rceil$.*

(iv) The conical hull of a point set S is defined as $\mathrm{cone}(S) = \left\{\sum_{p \in S} \alpha_p p : \alpha_p > 0 \text{ for all } p \in S\right\}$.

Proof: Let $h_1, \ldots, h_d$ be any d orthogonal halving hyperplanes of S and let the point p be the point common to $h_1, \ldots, h_d$. Consider any α-wedge whose apex is p and suppose that the axis of this wedge is the ray r. We claim that one of the planes h_i makes an angle of at least $\pi/2 - \arccos(1/\sqrt{d})$ with r. To see this, observe that the (positive and negative) normal vectors of the halving planes form a set of $2d$ points on the unit sphere in $\mathbb{R}^d$. In fact, they are the vertices of generalized octahedron. Placing spherical caps of angle $\arccos(1/\sqrt{d})$ gives a covering of the sphere, and hence r forms an angle of at most $\arccos(1/\sqrt{d})$ with at least one of the normals. Therefore, r forms an angle of at least $\pi/2 - \arccos(1/\sqrt{d})$ with the corresponding halving plane h_i.

Therefore, the α-wedge with axis r contains h_i and so contains one of the two halfspaces bounded by h_i. Since this is true for any α-wedge containing p we conclude that $D_\alpha(p, S) \geqslant \lceil n/2 \rceil$, as required. □

Theorem 2 *The function f_α^d satisfies*

$$\begin{array}{rcll} f_\alpha^d(n) & = & 1 & \text{if } \alpha < \pi \\ f_\alpha^d(n) & = & \lceil n/(d+1) \rceil & \text{if } \pi \leqslant \alpha < 2(\pi - \arccos(1/d)) \\ \lceil n/2 \rceil \geqslant f_\alpha^d(n) & \geqslant & \lceil n/(d+1) \rceil & \text{if } 2(\pi - \arccos(1/d)) \leqslant \alpha \leqslant \pi + 2\arccos(1/\sqrt{d}) \\ f_\alpha^d(n) & = & \lceil n/2 \rceil & \text{if } \pi + 2\arccos(1/\sqrt{d}) \leqslant \alpha < 2\pi \end{array}$$

Notice that, as $d \to \infty$, $2(\pi - \arccos(1/d)) \to \pi$ and $\pi + \arccos(1/\sqrt{d}) \to 2\pi$. Thus, Theorem 2 leaves a considerable gap in our knowledge.

4 Some Results for $\mathbb{R}^3$

Since we have been unable to fully determine f_α^d for all values of d, we concentrate our efforts in this section on the special case $d = 3$. We begin by restating Theorem 2 with $d = 3$.

Corollary 1 *The function f_α^3 satisfies*

$$\begin{array}{rcll} f_\alpha^3(n) & = & 1 & \text{if } \alpha < \pi \ (\alpha < 180^\circ) \\ f_\alpha^3(n) & = & \lceil n/4 \rceil & \text{if } \pi \leqslant \alpha < 2(\pi - \arccos(1/3)) \\ \lceil n/2 \rceil \geqslant f_\alpha^3(n) & \geqslant & \lceil n/4 \rceil & \text{if } 2(\pi - \arccos(1/3)) \leqslant \alpha < \pi + 2\arccos(1/\sqrt{3}) \\ f_\alpha^3(n) & = & \lceil n/2 \rceil & \text{if } \pi + 2\arccos(1/\sqrt{3}) \leqslant \alpha \leqslant 2\pi \end{array}$$

We first show that the situation is more complex in $\mathbb{R}^3$ than in $\mathbb{R}^2$. That is, the function f_α^3 does not change immediately from $\lceil n/4 \rceil$ to $\lceil n/2 \rceil$ at the threshold value $\alpha = 2(\pi - \arccos(1/3))$.

Lemma 10 *If $\alpha < 2(\pi - \arccos(1/\sqrt{5}))$ ($\alpha < 233.13^\circ$) then $f_\alpha^3(n) \leqslant 2\lceil n/5 \rceil$.*

Proof: Hardin, Sloane, and Smith [9] describe a covering of the unit sphere by 5 spherical caps whose angular radius[(v)] is $\arccos(1/\sqrt{5})$. Let the centers of these 5 caps be denoted by $v_1, \ldots, v_5$. (These are the vertices of a regular triangular bipyramid.) For the lower bound point set, we place $\lceil n/5 \rceil$ or $\lfloor n/5 \rfloor$, as appropriate, points on each of the 5 rays from the origin through $v_1, \ldots, v_5$, to produce set of n points.

The convex hull of $v_1, \ldots, v_5$ has 6 equilateral triangular faces. In particular, none of the faces are obtuse. Thus, in each face there is a point whose radial projection onto the unit sphere is contained in

(v) The *angular radius* of a point set S is $\max\{\angle pOq : p, q \in S\}$ where O is the origin.

three of the spherical caps. Stated another way, for each face f, there is a $2\arccos(1/\sqrt{5})$-wedge with apex at the origin that contains f.

At this point, the remainder of the proof is exactly as in the proof of Lemma 8. Take the conical hull of each face of the convex hull, and determine some conical hull h that contains $-p$. Then there is a $2\arccos(1/\sqrt{5})$-wedge with apex at p and that contains h so therefore contains at least $3\lfloor n/5\rfloor$ points of S. The complementary $(2\pi - 2\arccos(1/\sqrt{5}))$-wedge contains an α-wedge that contains p and at most $n - 3\lfloor n/5\rfloor \leqslant 2\lceil n/5\rceil$ points of S so $D_\alpha(p,S) \leqslant 2\lceil n/5\rceil$, as required. □

Next we give an improvement on the value of α required to achieve $f^3_\alpha(n) \geqslant \lceil n/2\rceil$.

Lemma 11 *If $\alpha \geqslant 3\pi/2$ ($\alpha \geqslant 270°$) then $f^3_\alpha(n) \geqslant \lceil n/2\rceil$.*

Proof: Fekete and Meijer [8] show that, for every set S of n points in $\mathbb{R}^3$, there exists 3 mutually orthogonal halving planes of S that partition $\mathbb{R}^3$ into 8 octants such that the number of points in opposite octants is equal. If we take p to be the common intersection point of these 3 halving planes and let w be any α-wedge whose apex is p then we find that, if some octant Q is not entirely contained in w then the octant $-Q$ is entirely contained in w. (This is because for any $q \in Q$ and $r \in -Q$, $\angle qOr \geqslant \pi/2$.) Thus, for every point of S not in w there is a point of S that is in w, so $D_\alpha(p,S) \geqslant \lceil n/2\rceil$, as required. □

The following theorem summarizes the above results.

Theorem 3 *The function f^3_α satisfies*

$$\begin{array}{rcll} f^3_\alpha(n) & = & 1 & \text{if } \alpha < 180° \\ f^3_\alpha(n) & = & \lceil n/4\rceil & \text{if } \pi \leqslant \alpha < 2(\pi - \arccos(1/3)) \\ \lceil 2n/5\rceil \geqslant f^3_\alpha(n) & \geqslant & \lceil n/4\rceil & \text{if } 2(\pi - \arccos(1/3)) \leqslant \alpha < 2(\pi - \arccos(1/\sqrt{5})) \\ \lceil n/2\rceil \geqslant f^3_\alpha(n) & \geqslant & \lceil n/4\rceil & \text{if } 2(\pi - \arccos(1/\sqrt{5})) \leqslant \alpha < 3\pi/2 \\ f^3_\alpha(n) & = & \lceil n/2\rceil & \text{if } 3\pi/2 \leqslant \alpha < 2\pi \end{array}$$

5 Conclusions

We have completely determined the function f^2_α and given a linear-time algorithm for finding a point p such that $D_\alpha(p,S) \geqslant f^2_\alpha(|S|)$. Our main new algorithmic result is a linear-time algorithm for finding 3 concurrent halving lines, each pair of which forms an angle of $\pi/3$. These triples of halving lines were used by Fekete and Meijer [8] to show that the cost of a minimum Steiner star of an n point set in $\mathbb{R}^2$ is at most $2/\sqrt{3}$ times the cost of the maximum matching of the same set. Our algorithm gives an $O(n)$ time construction of a Steiner star matching this bound.

Fekete and Meijer also prove that, in $\mathbb{R}^3$, the ratio between the minimum Steiner star and the maximum matching is at most $\sqrt{2}$ by showing the existence of 3 orthogonal halving planes with the property that the number of points in opposite orthants is equal. They prove this by taking an arbitrary halving plane Π, projecting the points onto Π, and finding two orthogonal halving lines in Π such that opposite quadrants have the same number of projected points above and below Π. The existence of these two halving lines is guaranteed by a simple continuity argument. A (simpler) variant of the algorithm from Lemma 3 can be used to find these two orthogonal halving lines and hence find three orthogonal halving planes in $O(n)$ time. Again, this gives an $O(n)$ time algorithm to construct the Steiner star achieving this bound.

We conclude with a list of open problems:

1. Given a point set S in $\mathbb{R}^2$, what is the complexity of finding a point $p \in \mathbb{R}^2$ that maximizes $D_\alpha(p, S)$? For $\alpha = \pi$, i.e., halfplane depth, Chan has recently given an $O(n \log n)$ time algorithm [4].

2. Our understanding of the function f_α^d is still incomplete for $d \geqslant 3$.

3. In $\mathbb{R}^2$, we are able to find 3 concurrent halving lines whose sides are parallel to the edges of an equilateral triangle. The same technique can be used in $\mathbb{R}^d$, if d is even, to show that there always exists $d + 1$ concurrent halving hyperplanes whose sides are parallel to the edges of a regular d-simplex. (The proof involves continuously rotating the d-simplex until each of its vertices has been reflected through the origin; this works in even dimensions because reflection through the origin can be implemented as a sequence of rotations.) A different proof can be used to prove the same result for $\mathbb{R}^3$. Does this result hold in $\mathbb{R}^d$ for all values of d?

4. In $\mathbb{R}^2$ any pair of orthogonal halving lines partitions the plane into four quadrants such that the number of points in opposite quadrants is equal. In $\mathbb{R}^3$, Fekete and Meijer [8] showed the existence of 3 orthogonal halving planes with the same property. This raises the following question: Given a set S of n points in $\mathbb{R}^d$ is it the case that we can always find an arrangement of d mutually orthogonal halving hyperplanes of S such that cells with opposite sign vectors in the arrangement contain the same number of points of S?

5. Given a set S of n points in $\mathbb{R}^d$, for $d \geqslant 3$, what is the complexity of finding a point $p \in \mathbb{R}^d$ such that $p \geqslant f_\alpha^d(n)$? This problem is still open even for the case $\alpha = \pi$, though the algorithm of Jadhav and Mukhopadhyay [10] settles the problem for $d = 2$ and $\alpha = \pi$.

6. In this paper we have only considered α-wedges. This is mainly because, for $d = 2$, these are more or less the only interesting scale-invariant objects. However, in higher dimensions, one can define many scale-invariant shapes. In general, for any shape F, one can study the properties of F-depth:

$$D_F(p, S) = \min\{h \cap S : h \text{ is an } F \text{ that contains } p\} \quad .$$

Acknowledgements

This research was initiated while Pat Morin was visiting Universitat Politècnica de Catalunya (UPC). The authors would like to thank Jit Bose, Anil Maheshwari, Perouz Taslakian, Michiel Smid, and David Wood for helpful discussions.

References

[1] M. Abellenas, M. Claverol, F. Hurtado, and C. Seara. (a, k)-sets en el plano. In *Proc. III Jornadas Espaolas de Matemtica Discreta*, pages 11–14, 2002.

[2] D. Avis, B. Beresford, L. Devroye, Hossam. Elgindy, Eric Guévremont, F. Hurtado, and B. Zhu. Unoriented θ-maxima. complexity and algorithms. *SIAM Journal on Computing*, 28(1):278–296, 1999.

[3] M. Blum, R. W. Floyd, V. Pratt, R. R. Rivest, and R. E. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7(4):448–461, 1973.

[4] T. M. Chan. An optimal randomized algorithm for maximum Tukey depth. In *Proc. 15th ACM-SIAM Symposium on Discrete Algorithms (SODA 2004)*, pages 423–429, 2004.

[5] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Heidelberg, 1997.

[6] J. Eckhoff. Helly, Radon, and Carathodory type theorems. In P. M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*, chapter 2.1, pages 389–448. North-Holland, Amsterdam, Netherlands, 1993.

[7] V. Estivill-Castro, Joseph O'Rourke, J. Urrutia, and D. Xu. Illumination of polygons with vertex floodlights. *Information Processing Letters*, 56:9–13, 1995.

[8] S. P. Fekete and H. Meijer. On minimum stars and maximum matchings. *Discrete and Computational Geometry*, 23:389–407, 2000.

[9] H. Hardin, N. J. A. Sloane, and W. D. Smith. Spherical codes. Book in preparation.

[10] S. Jadhav and A. Mukhopadhyay. Computing a centerpoint of a finite planar set of points in linear time. *Discrete & Computational Geometry*, 12:291–312, 1994.

[11] C.-Y. Lo, J. Matousek, and W. Steiger. Algorithms for ham-sandwich cuts. *Discrete & Computational Geometry*, 11:433–452, 1994.

[12] J. Matoušek. Approximations and optimal geometric divide-and-conquer. In *Proc. 23rd Ann. ACM Symposium on Theory of Computing*, pages 505–511, 1991.

[13] J. S. B. Mitchell and Joseph O'Rourke. Computational geometry column 42. *International Journal of Computational Geometry and Applications*, 11(5):573–582, 2001.

[14] J. Tukey. Mathematics and the picturing of data. In *Proc. Int. Congress of Mathematicians*, volume 2, pages 523–531, 1975.