Paths with no small angles

IMRE BÁRÁNY

Rényi Institute of Mathematics
Hungarian Academy of Sciences
POBox 127, 1364 Budapest, Hungary
e-mail: barany@renyi.hu and
Department of Mathematics
University College London
Gower Street, London WC1E 6BT
England

ATTILA PÓR

Department of Applied Mathematics, Charles University,
Malostranské nám. 25, 118 00 Praha 1, Czech Republic
e-mail: apor@renyi.hu

PAVEL VALTR

Department of Applied Mathematics, Charles University
Malostranské nám. 25, 118 00 Praha 1, Czech Republic
http://kam.mff.cuni.cz/∼valtr/

March 14, 2009

Abstract

Giving a (partial) solution to a problem of S. Fekete [3] and S. Fekete and G.J. Woeginger [4] we show that given a finite set X of points in the plane, it is possible to find a polygonal path with $|X|-1$ segments and with vertex set X so that every angle on the polygonal path is at least $\pi/9$. According to a conjecture of Fekete and Woeginger, $\pi/9$ can be replaced by $\pi/6$. Previously, the result has not been known with any positive constant. We show further that the same result holds, with an angle smaller than $\pi/9$, in higher dimensions.

1 Introduction and results

1.1 The plane

The aim of this paper is to answer the following beautiful and inspiring question which appeared first in S. Fekete's thesis in [3] in 1992, and later in the paper by Fekete and Woeginger [4] in 1997. The question is this. Given a finite set X of points in the plane, is it possible to find a polygonal path with $|X| - 1$ segments and with vertex set X so that every angle on the path is at least α (for some universal constant $\alpha > 0$)? The answer is, as we shall see soon, yes. This might be a first step toward proving a conjecture of S. Fekete and G.J. Woeginger [3, 4] that this result holds with $\alpha = \pi/6$. We prove the result with the constant $\alpha = \pi/9$. First we introduce notation and terminology.

Let $A_0, A_1, \ldots, A_n$ be $n+1$ distinct points in the plane (or, more generally, in d-dimensional space). We denote the path consisting of the segments A_0A_1, A_1A_2, $\ldots$, $A_{n-1}A_n$ by $A_0A_1 \ldots A_n$. This is a polygonal path with vertices $A_0, A_1, \ldots, A_n$. The angle of this path at A_i is the angle of the triangle $A_{i-1}A_iA_{i+1}$ at vertex A_i, $1 \leq i < n$.

Definition. Let $\alpha > 0$. We call the path $A_0A_1 \ldots A_n$ α-*good* if the angle at A_i is at least α for every $1 \leq i < n$. A path in the plane is called *good*, if it is $\pi/9$-good.

The main result of this paper is the following

Theorem 1. *For every finite set of points X in the plane there exists a good path on the points of X (containing each point of X as a vertex exactly once).*

We mention that $\pi/9$ in the theorem cannot be replaced by anything larger than $\pi/6$. This is shown when X consists of the center and the three vertices of a regular triangle (see Figure 1) when $|X| = 4$. This can be extended to arbitrarily large (even infinite) $|X|$ by placing a small copy of the 4-point example near the origin, and adding points of the form $(k, 0)$ to X where k is an integer.

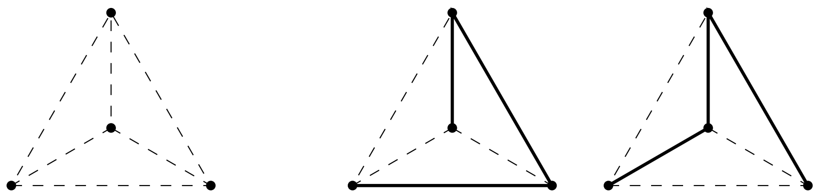

Figure 1: A 4-point configuration and its two paths

Another example, depicted in Figure 2, shows that Theorem 1 cannot be strengthened to paths with no self-intersections. It also shows that paths minimizing various quantities (such as total length, total turning angle) may have an arbitrarily small angle.

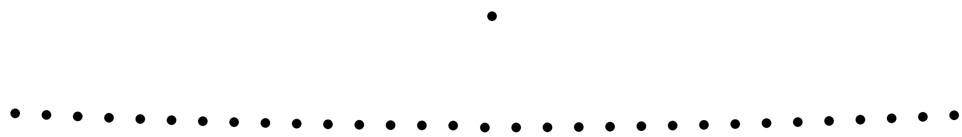

Figure 2: Every good path on this point set is self–intersecting (the set consists of points on a huge circle and one extra point inside the circle)

We will prove a slightly stronger statement which is more convenient for the induction argument. We will need two additional definitions.

Definition. We call the (oriented) directions of the vectors $\overline{A_1A_0}$ and $\overline{A_{n-1}A_n}$ the two *end directions* of the path $A_0 \dots A_n$. We identify the (oriented) directions with points of the unit circle S^1.

In the following definition and in the proof of Theorem 1 we fix $\alpha = \pi/9$.

Definition. We call a subset R of the unit circle a *restriction* if it is the disjoint union of two intervals $R_1, R_2 \subset S^1$ such that both have length $4\alpha = 4\pi/9$ and their distance from each other (along the unit circle) is at least $2\alpha = 2\pi/9$. We call the path $A_0 \dots A_n$ *R-avoiding* if the two end directions are not in the same R_i ($i = 1, 2$) and the path is good (see Figure 3).

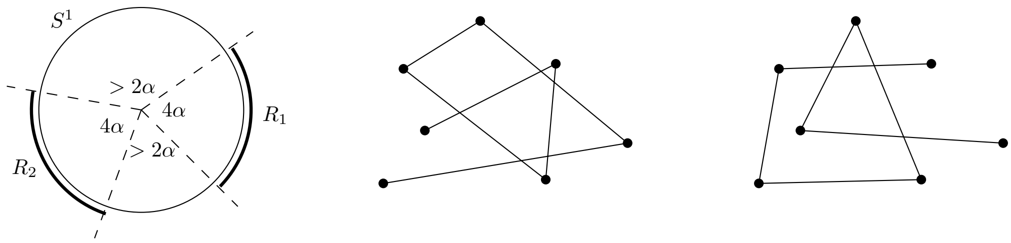

Figure 3: A restriction $R = R_1 \cup R_2$ and two (good) paths that are not R-avoiding.

The following theorem is a strengthening of Theorem 1.

Theorem 2. *Let X be a finite set of points in the plane. For every restriction R there is an R-avoiding path on all the points of X.*

The proof of this theorem goes by induction on $n = |X|$, giving a straightforward $O(n^2 \log n)$ algorithm for finding a $\pi/9$-good path. The running time can be improved to $O(n^2)$, when one uses the convex hull algorithm of [5], say. A sketch of an $O(n^2)$ algorithm can be found in the conference version [1] of this paper.

1.2 Higher dimensions

The natural question is what happens in higher dimensions. In the final section of this paper we prove the following result.

Theorem 3. *There is a positive α such that for every $d \geq 2$ and for every finite set of points X in d-dimensional space there exists an α-good path on X.*

Actually, the proof method of Theorem 2 works but some extra difficulties have to be overcome. We get $\alpha = \pi/42$ from the proof. Perhaps the example in Figure 1 is the extremal case in all dimensions:

Conjecture. Theorem 3 holds with $\alpha = \pi/6$.

1.3 An open problem

Another problem that we encountered while working on this paper seems interesting and nontrivial. Call a finite set X in the d-dimensional space α-*flat* if every triangle with vertices from X has an angle smaller than α. One example of an α-flat set is a finite set X_0 of collinear points. Each point of X_0 can be moved freely in a small enough neighbourhood so that the resulting set X_1 is still α-flat. Next, each point of X_1 can be replaced by a very small but otherwise arbitrary α-flat set, and the resulting set is still α-flat if the replacements are small enough. Perhaps all α-flat sets can be obtained by repeating this process a finite number of times.

Next, call the set X β-*separable* if it can be partitioned as $X = U \cup V$ with U, V disjoint and nonempty so that the angle between the line through u_1, v_1 and the line through u_2, v_2 is smaller than β for every $u_1, u_2 \in U$ and every $v_1, v_2 \in V$.

Conjecture. For every $d \geq 2$ and for every positive β there is a positive $\alpha_d(\beta)$ such that every α-flat set X in d-dimensional space is β-separable.

We have a proof of this conjecture for $d = 2$.

2 Proof of Theorem 2

We prove the theorem by induction on the number of points in X. In this section we fix α as $\pi/9$.

If $|X| = 2$ then the two end directions are the opposite to each other. Since the length of R_i, $4\alpha = 4\pi/9$, is smaller than π the two end directions cannot be in the same interval R_i.

Assume $|X| > 2$. Let K be the convex hull of X and let $V \subseteq X$ be the vertex set of K. Next let $R = R_1 \cup R_2$ be a restriction. We distinguish two cases depending on the smallest angle of the polygon K.

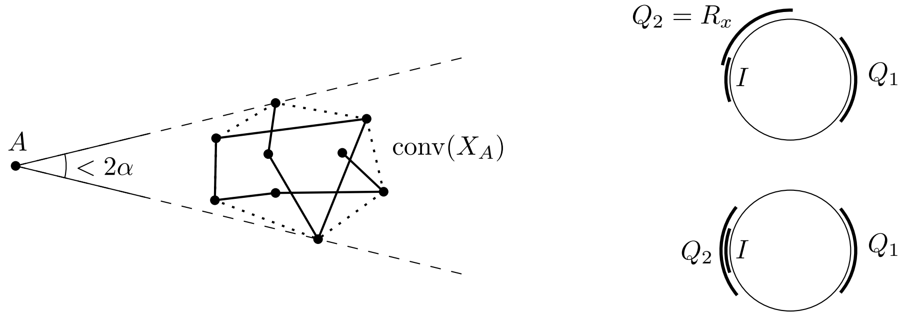

Figure 4: Case 1

Case 1: The smallest angle of K is smaller than 2α. Let A be the vertex where that smallest angle occurs and let $X_A = X \setminus \{A\}$. We can assume, without loss of generality, that X_A is contained in the wedge of angle 2α whose vertex is A and whose line of symmetry is the x-axis, see Figure 4. Then for each point $B \in X_A$ the direction $\overline{BA}$ is in the interval $I = (\pi - \alpha, \pi + \alpha) \subset S^1$. Since the length of I is 2α it can only intersect one of the two intervals R_1 and R_2. Let $Q_1 = [-2\alpha, 2\alpha] \subset S^1$. If one of the sets R_1 or R_2 intersects I, then let Q_2 be equal to the one that intersects I. Otherwise set $Q_2 = [\pi - 2\alpha, \pi + 2\alpha]$. It is easy to see that $Q = Q_1 \cup Q_2$ is a restriction; this is the point where the bound $\alpha \leq 20^0 = \frac{\pi}{9}$ comes from. By induction we find a Q-avoiding path $p = A_0A_1 \ldots A_n$ on X_A. If the end direction $\overline{A_1A_0}$ is not in Q_1, then we can extend this path to the good path $Ap = AA_0 \ldots A_n$ on X. Analogously, if the end direction $\overline{A_{n-1}A_n}$ is not in Q_1, then we can extend this path to the good path $pA = A_0 \ldots A_nA$ on X.

So at least one of the extended paths pA, Ap is α-good. The end direction at A is always in I. Therefore, if both end directions of Ap or of pA are in R_1 (or R_2), then both have to be in Q_2. In this case we can extend p at both ends. But only one of the end directions of p is in Q_2. So we extend p at the end which is in Q_2 and we get an R-avoiding path on X.

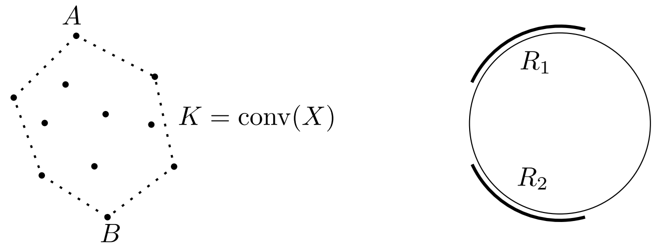

Figure 5: Case 2

Case 2: Every angle of K is at least 2α. See Figure 5. Without loss of generality we can assume that R_1 and R_2 are symmetric to the horizontal line. Let A and B be the vertices of K with the largest and smallest y-coordinate, respectively. We will distinguish three subcases depending on the size of $Y = X \setminus V$.

Case 2a: The set Y is empty. As a first attempt we try to find an R-avoiding path that contains only edges of K. Such a path can be identified by the missing edge of K. All these paths are clearly α-good. If there is an edge on the perimeter of K with a direction not in R_1 or R_2, then the path missing the next edge will have that direction as end direction. In this case we have found an R-avoiding path.

Now we assume that for each edge in K one direction is in R_1 and the other in R_2. If $|X| > 4$, then there is a path along the perimeter of K between A and B of length at least three. Take the path that misses an edge disjoint from A and B — see Figure 6 (left). One of the end directions will be in the interval $[0, \pi]$ (upwards) and the other one will be in $[\pi, 2\pi]$ (downwards). This path is R-avoiding since $R_1 \subset (0, \pi)$ and $R_2 \subset (\pi, 2\pi)$.

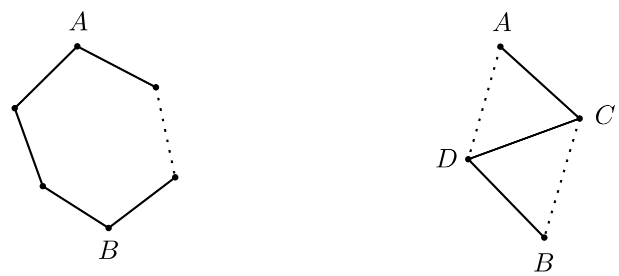

Figure 6: Case 2a

If $|X| = 3$ then the path missing edge AB from K is R-avoiding since it has one upward and one downward end direction.

If $|X| = 4$ and AB is an edge of the convex hull, then the path missing this edge is R-avoiding. If A and B are opposite vertices of K (which is a quadrilateral now), then we connect the four vertices from top to bottom starting with A and ending with B. Let this path be $ACDB$ — see Figure 6 (right). We have CA in R_1 and CD is pointing downwards. That is $\overline{CA} \in [\alpha, \pi - \alpha]$ and $\overline{CD} \in [\pi, 2\pi]$ and therefore the angle at C is at least α. Similarly the angle at D is at least α as well which shows that this path is good. The end directions are again upward and downward therefore the path $ACDB$ is an R-avoiding path.

Case 2b. The set Y consists of one point: $Y = \{F\}$, say. Take a path that contains all edges of K except one and the segment from F to one of the endpoints of the missing edge. If the angle at the vertex which is connected to F is at least α we have a good path.

In this way every segment from F to a vertex of K can be extended to a good path since each angle of K is at least 2α and therefore the angle toward one of the neighbours along the perimeter of K has to be at least α.

Consider the extended path starting with FB — see Figure 7 (left). The end direction $\overline{BF}$ is upwards. If $\overline{BF}$ or the other end direction is not in R_1 we have an R-avoiding path.

Figure 7: Case 2b

If the other end direction is in R_1, then it directs upwards which can only occur if AB is an edge of the convex hull and the path extended from FB ends at A.

Similarly the path extended from FA will end in B so we found an α-good Hamiltonian cycle — see Figure 7 (right). If X has at least five elements, then there is an edge of K which is disjoint from A and B and we can use a previous argument. If X has four elements, then we take the path going from top to bottom starting at A and ending at B. In both cases the arguments are identical to the ones in **Case 2a**.

Case 2c. The set Y has at least two elements. By induction we find an R-avoiding path $p = A_0 \ldots A_n$ on Y. We will extend this path as follows. Let $F \in V$, that is, F a vertex of K. Connect A_0 (resp. A_n) to F and then connect F to one of its neighbours, G say, on the convex hull. Continue the path along the convex hull, we get a new path p^*. This path can be written as $p^* = ..GFp$ or $pFG..$, where the two dots represent the unique continuation of the path along the perimeter of K. The path p^* will be good if the angles at A_0 (resp. A_n) and at F are at least α.

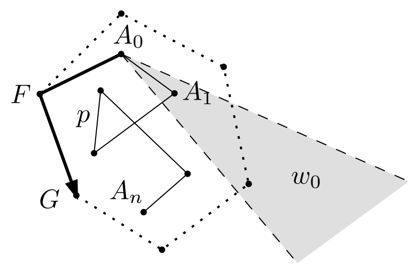

Figure 8: The wedge w_o

Consider first the angle at A_0 (resp. A_n). Let w_0 be the set of all points W for which the angle A_1A_0W is smaller than α — see Figure 8. Similarly let w_n be the set of all points W for which the angle $A_{n-1}A_nW$ is smaller than α. Both sets w_0 and w_n are wedges with an angle of 2α. The angle of p^* at A_0 (resp. A_n) is at least α if and only if F is not in the wedge w_0 (resp. w_n). Observe that V is not contained in w_0 as otherwise A_0 would be a vertex of K. Thus we can choose $F \in V$ so that the angle at A_0 is at least α — see Figure 8. Similarly, V is not contained in w_n, and we can choose F so that the angle at A_n is at least α.

Consider now the angle at F. To continue the path from F we have two choices for G to go along the perimeter of K. Since the angle at each vertex of K is at least 2α, one of the choices certainly yields a path whose angle at F is at least α. Consequently there is at least one good path p^* of the form $..GFp$ and one of the form $pFG..$ (the two Gs may be distinct).

One end of such a p^* is an edge of K and the other one is A_1A_0 or $A_{n-1}A_n$. If $\overline{A_1A_0}$ or $\overline{A_{n-1}A_n}$ is not in R, then we keep the end which is not in R and extend the path through the other end to get a good path on X which will be R-avoiding.

Thus we can assume that $\overline{A_1A_0}$ is in R_1 and $\overline{A_{n-1}A_n}$ is in R_2, say. This has the beneficial consequence that A is not in w_0 as the wedge w_0 lies completely below the horizontal line trough A, further denoted by l — see

Figure 9 (left). Thus A can be taken for F and there is a good path of the form $p^* = ..GAp$. Similarly, $B \notin w_n$ and there is a good path $p^* = pBG...$

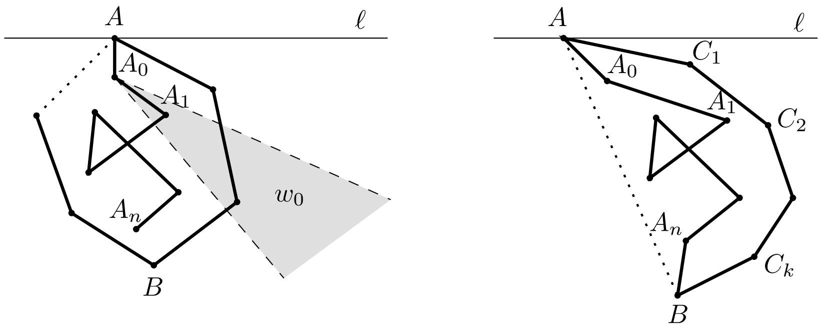

Figure 9: Case 2c

Notice now that $p^* = ..GAp$ is R-avoiding unless both of its end directions are in R_2. This can only happen if AB is an edge of K and the angle A_0AB is smaller than α. Similarly, $p^* = pBG..$ is R-avoiding unless both of its end directions are in R_1. This can only happen if AB is an edge of K and the angle A_nBA is smaller than α.

In this situation let $A, C_1 \dots, C_k, B, A$ be the vertices of K in this order. It follows that all angles along the Hamiltonian cycle

$$A, C_1, \dots, C_k, B, A_n, A_{n-1} \dots, A_0, A$$

are at least α. See Figure 9 (right). As we have seen in **Case 2b**, such a cycle produces an R-avoiding path unless $k = 1$.

The only remaining case is when $k = 1$, then K is the triangle ABC where we set $C = C_1$. Observe now that $|V \cap w_0| \leq 1$, since the angle at A of K is at least 2α and so w_0 cannot contain both B and C. Similarly, $|V \cap w_n| \leq 1$.

We assume next that the angle A_1A_0C is at least α, that is $C \notin w_0$. If the angle A_0CB is at least α, then the path $A_n \dots A_0CBA$ is R-avoiding — see Figure 10 (left). Otherwise the angle A_0CA is at least α and the path $BA_n \dots A_0CA$ is R-avoiding. From now on we can assume that $C \in w_0$.

Similarly we can find an R-avoiding path if the angle $A_{n-1}A_nC$ is at least α. From now on we can assume that $C \in w_n$.

This implies $V \cap w_0 = V \cap w_n = \{C\}$. Thus p can be extended to a good path p^* at both ends through both A and B.

The angle A_nAC has to be smaller than α as otherwise $A_0 \dots A_nACB$ is an R-avoiding path. Similarly the angle A_0BC is smaller than α. We have seen above that

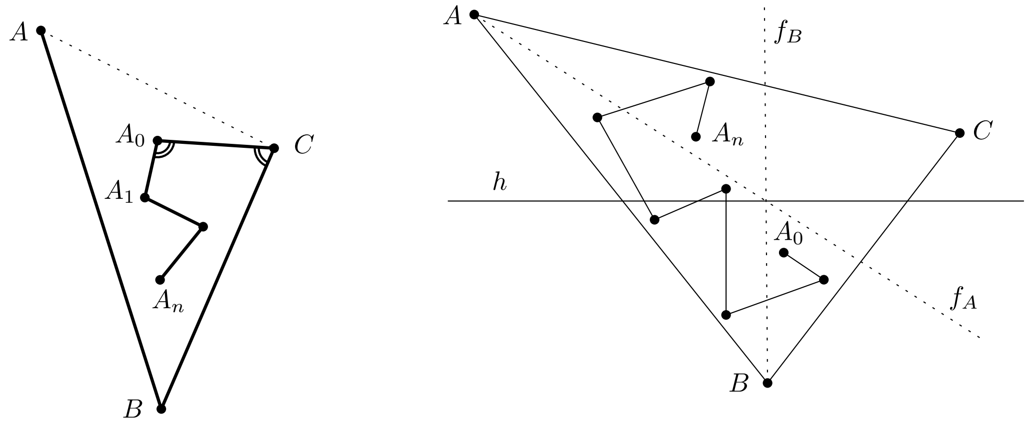

Figure 10: Case 2c when the convex hull is a triangle

$\angle A_0AB < \alpha$ and $\angle A_nBA < \alpha$.

Now let f_A (resp. f_B) be line through A (and B) halving the angle BAC (and the angle ABC). See Figure 10 (right). Let h be the horizontal line through the intersection of f_A and f_B. What we established so far implies that A_0 (resp. A_n) is in the triangle delimited by f_A, f_B, BC and by f_A, f_B, AC.

It follows then that A_0 is below and A_n is above h. Now w_0 lies entirely below h and w_n lies entirely above h, contradicting $C \in w_0 \cap w_n$.

3 Higher dimensions

Throughout this section we consider α very small, say $\alpha = 0.1^o$. We do so in order to simplify the computations. Actually, the proof below gives $\alpha = \pi/42 = 4.2857\ldots^o$, when the computations are done properly. We mention without proof that a more complicated argument gives a somewhat bigger α.

We identify the unit sphere S^{d-1} with the set of (oriented) directions in the d-dimensional space. A subset R of the unit sphere S^{d-1} is called a *restriction* if it is the disjoint union of two spherical caps, R_1 and R_2, each of (spherical) diameter 10α such that the (spherical) distance of R_1 and R_2 is at least 8α. More precisely, each R_i is a set of directions differing from a fixed direction by at most 5α, and each direction in R_1 differs from each direction in R_2 by at least 8α. Again, a path is *R-avoiding* if it is α-good and its two end directions are not in the same R_i. If a path is α-good then we say shortly that it is *good*.

Theorem 4. *Let X be a finite set of points in some Euclidean space (of dimension d). For every restriction R there is an R-avoiding path on all the points of X.*

Proof: We proceed by induction on $|X|$. Proving the starting steps ($|X| < 6$) of the induction is tiresome and not quite simple. We postpone it to the next section because it uses the proof scheme of the general induction step which follows now.

So we assume that $|X| \geq 6$. Let $A, B \in X$ be two points of X such that AB is a diameter of X. We will distinguish three cases.

Case 1. For any point $P \in X$ different from A, B the angle $\angle BAP \leq 4\alpha$. Or analogously, $\angle ABP \leq 4\alpha$ for all $P \in X$, different from A, B.

In this case we basically repeat the proof of Theorem 2 in Case 1. For a direction $d \in S^{d-1}$ and for an angle $\phi \in (0, \pi)$, we denote by $C(d; \phi)$ the cap of S^{d-1} consisting of (oriented) directions differing from d by at most ϕ. The roles of the intervals I and Q_1 are now played by the caps $C(\overline{BA}; 4\alpha)$ and $C(\overline{AB}; 5\alpha)$, respectively, and the role of Q_2 is now played either by an R_i intersecting $C(\overline{BA}; 4\alpha)$ (if such an R_i exists, in which case it is unique) or by the cap $C(\overline{BA}; 5\alpha)$ (otherwise). We remark that we now need $\alpha \leq \pi/27$ to make sure that $Q_1 \cup Q_2$ is a restriction. Then we may use exactly the same arguments as in the plane.

Case 2. We find two points $C, D \in X$ such that the following hold (see Figure 11). The angles $\angle DAB, \angle ABC, \angle DAC$ and $\angle DBC$ are at least 2α. Further, the angles $\angle BCD$ and $\angle CDA$ are at least α. (Note that $\angle BDA, \angle BCA \geq 60^o > \alpha$ since AB is the diameter of X.)

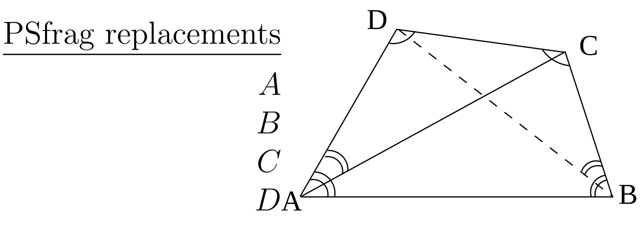

Figure 11: Case 2

This case is fairly straightforward. First we find an R-avoiding path p on $X \setminus \{A, B, C, D\}$. The argument from Section 2 shows that either pA or pB is a good path. We assume without loss of generality that pA is a good path. If we can continue it toward D, then both $pADBC$ and $pADCB$ are full extensions. Obviously one of them is R-avoiding. If pA does not extend toward D, then both $pABCD$ and $pACBD$ are full extensions. One of them is R-avoiding unless $\overline{CD}$, $\overline{BD}$, and the first end direction of p lie in the same R_i, say R_1.

If both ends of p extend to A, then the same arguments apply. We conclude that both $pABCD$ and $DCBAp$ are good paths. One of them is clearly R-avoiding (we use that the end directions of p cannot lie in the same R_i).

Thus p cannot be extended to A at both ends implying that Bp is a good path. The same arguments apply again showing that $\overline{DC}$, $\overline{AC}$, and the last end direction of p all lie in R_2. We observe, finally, that $DBpAC$ is a good path which is R-avoiding as well since $\overline{BD} \in R_1$ and $\overline{AC} \in R_2$.

Case 3 When the conditions of Case 1 and 2 fail to hold.

First we show that there exists a point F in X such that the angles $\angle FAB$ and $\angle ABF$ are both at least 2α.

Since we are not in Case 1 we have a point C such that $\angle ABC \geq 4\alpha$ and a point D such that $\angle BAD \geq 4\alpha$. If the two points C and D coincide, then this point will do for F. If the angle $\angle BAC$ or $\angle ABD$ is at least 2α, then C or D will do as F. Otherwise $\angle BAD, \angle CAD, \angle ABC, \angle CBD$ are all at least 2α. A little elementary 3-dimensional calculation (we omit the details) shows that $\angle ADC, \angle DCB \geq \alpha$ implying that C and D are two points satisfying the conditions of Case 2.

Let $p = ED \ldots D'E'$ be an R-avoiding path on $X \setminus \{A, B, F\}$. We can extend p at either end to A or B and then to a good path on X. Obviously one of them will be R-avoiding except when $\overline{DE}$ is in one of R_1 and R_2 and $\overline{D'E'}$ is in the other one. Assume (without loss of generality) that $\overline{DE} \in R_1$ and $\overline{D'E'} \in R_2$.

First we show how to find an R-avoiding path if one of the R_i, say R_1, has a direction closer than α to a direction perpendicular to AB. One of the paths pA or pB is good. We assume, again without loss of generality, that pA is good. One of the paths $pABF$ or $pAFB$ is certainly good and then it is R-avoiding except if $\overline{BF}$ or $\overline{FB}$ is in R_1. Therefore BF is *almost* perpendicular to AB, meaning that $\angle ABF > \pi/2 - 11\alpha$.

As p can be extended at the other end, one of the paths Ap or Bp is good. Assume first that Ap is good, then so is $FBAp$ or $BFAp$. Then one direction of the line BF is in R_1 and the other one is in R_2.

We claim that in this case Bp cannot be a good path. If it were, then its full extension would have an end direction in the line AF. We may suppose that this end direction is in R_1 or in R_2. Now AB is a diameter of X so $\angle AFB \geq \angle ABF > \pi/2 - 11\alpha$. On the other hand, $\angle AFB < \pi - \angle ABF < \pi/2 + 11\alpha$. Thus each of the two directions of line AF differs from each of the directions of line BF by more than $\pi/2 - 11\alpha \geq 10\alpha$, which is a contradiction, since R_1 or R_2 contains one direction of each of the lines AF, BF (note that this is the place where we needed $\alpha \leq \pi/42$). This proves our claim and shows, further, that $\angle DEB < \alpha$, implying further that the directions $\overline{DE}$ and $\overline{BE}$ differ by at most α.

We have to consider two simple cases now. We write cone(P, UV, γ) for the circular cone with apex P, axis going in direction $\overline{UV}$, and half-angle γ.

Case a. $\overline{FB} \in R_1$ and $\overline{BF} \in R_2$. Then $BFAp$ is not good so $\angle EAF < \alpha$. So $E \in$ cone(A, AF, α). Both $\overline{DE}$ and $\overline{FB}$ lie in R_1. So direction $\overline{BE}$ differing from $\overline{DE}$ by at most α differs from $\overline{FB}$ by at most 11α, implying $E \in$ cone$(B, FB, 11\alpha)$. This is impossible: the two cones have no point in common since $\angle BAF > 2\alpha$, see Figure 12 where $\beta = 11\alpha$.

PSfrag replacements
A
B
F
α
β

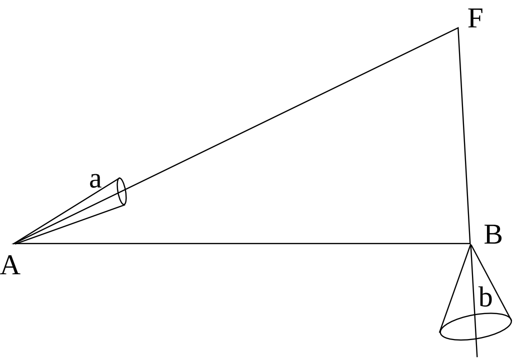

Figure 12: The two cones have no common point

Case b. $\overline{BF} \in R_1$ and $\overline{FB} \in R_2$. Then $FBAp$ is not good, thus $\angle EAB < \alpha$ and so $E \in$ cone(A, AB, α). Also, $E \in$ cone$(B, BF, 11\alpha)$ as $\overline{BF}, \overline{DE} \in R_1$ and $\overline{DE}$ and $\overline{BE}$ differ by at most α. It is easy to check that in this case Fp is a good path, which has a full extension since the angle at F is large. But this extension is R-avoiding since one of its end directions is contained in the line AB which is almost perpendicular to both R_1 and R_2, see Figure 13 where $\beta = 11\alpha$, again.

PSfrag replacements
A
B
F
α
β

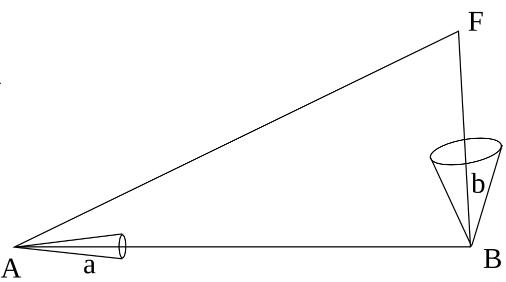

Figure 13: The two cones intersect

We are finished with the case when Ap is a good path. Assume now that Ap is not a good path. Then $\angle DEA < \alpha$ and therefore AE is almost perpendicular to AB. By elementary geometry we get that E and F satisfy the conditions of Case 2, which is a contradiction, again.

From now on we assume that both R_1 and R_2 are at distance at least α from the great sphere $h \subset S^{d-1}$ which is perpendicular to AB. The sphere S^{d-1} is cut by h into two halfspheres h_A and h_B where h_A contains the direction $\overline{BA}$. Obviously, each R_i is contained in h_A or h_B, and if $R_1 \subset h_A$, say, then Ap is a good path.

If both $R_1, R_2 \subset h_A$, then both ends of p can be extended to A and then to full extensions with opposite end directions $\overline{BF}$ and $\overline{FB}$, and at least one of them is neither in R_1 nor in R_2, and we have an R-avoiding path.

Suppose finally that $R_1 \subset h_A$ and $R_2 \subset h_B$. Both paths pB and Ap are good. If $FBAp$ (or $pBAF$) is a good path, then it is R-avoiding because $\overline{BF} \in h_A$ (or $\overline{AF} \in h_B$). Otherwise the Hamiltonian cycle $FApBF$ is good, that is, all the angles along the cycle are at least α.

Everything is under control now. Removing any edge from this cycle produces a good path. If none of these yield an R-avoiding path, then all edges of this cycle belong to R_1 in one direction, and to R_2 in the other. Thus, any two edges of this cycle are almost parallel. Moreover, going along one direction in this cycle the direction of the edges is as follows $R_1, R_1, R_2, R_2, R_1, R_1, R_2, R_2, \ldots$. Indeed, if there were consecutive R_i, R_j, R_i in the sequence of directions (with $i = j$ or not), then deleting the middle edge would produce an R-avoiding path.

Observe now that $\overline{AB} \in R_2$ since $\overline{AF}, \overline{FB} \in R_2$, and the vector AB is the sum of the vectors AF and FB. This shows that every direction in R_2 is closer than 10α to $\overline{AB}$. Similarly, every direction in R_1 is closer than 10α to $\overline{BA}$.

Assume now that $\overline{EF} \in h_A$. Then the path $AFpB$ is R-avoiding: the only angle to be checked is $\angle EFA$ but there $\overline{FA}$ is close to $\overline{BA}$ and the angle between directions $\overline{EF}$ and $\overline{BA}$ is at most $\pi/2$.

Thus, finally, $\overline{EF} \in R_2$. Set $p^* = p \setminus E$. We claim that the path p^*BFEA is R-avoiding. The only critical angle is $\angle FEA$ and here $\overline{AF}$ is close to $\overline{AB}$ and the angle between directions $\overline{FE}$ and $\overline{AB}$ is at most $\pi/2$.

4 Starting the induction

The case $|X| = 2$ is trivial. If $X = \{A, B, C\}$ and AB is a diameter of X then $\angle ACB \geq \pi/3 > 10\alpha$ and thus the path ACB is R-avoiding for any restriction R.

Consider next the case $|X| = 4$. Then X lies, of course, in 3-dimensional space. If we are in **Case 1** of the preceding section, then we need the induction basis for $X \setminus A$, which has three elements and that case has been covered. So assume $|X| = 4$ and we are not in **Case 1**. Then at most one angle is smaller than α at every vertex: this is clear at the endpoints of the diameter AB, and if at vertex C, say, both $\angle ACD$ and $\angle DCB$ are smaller than α, then $\angle ACB < 2\alpha$, yet $\angle ACB \geq \pi/3$ as AB is the diameter.

We assume now that R_1 and R_2 are symmetric with respect to a horizontal plane. Let T, U, V, Z be the points of X in vertically decreasing order. (We need new notation for the points, and we will only use the fact that at most one angle is smaller than α at every vertex.)

If the path $TUVZ$ is not R-avoiding, then $\angle TUV < \alpha$ or $\angle UVZ < \alpha$. Without loss of generality we can assume that $\angle TUV < \alpha$. Then, just as in **Case 2a** of the planar case, the line TU is almost horizontal implying that $\overline{UT} \notin R_1$ and $\overline{TU} \notin R_2$. Next, $TUZV$ is R-avoiding unless $\angle UZV < \alpha$. Then $VZTU$ is R-avoiding unless $\angle ZTU < \alpha$. But in this case the path $VTZU$ is R-avoiding.

Consider now the case $|X| = 5$. If we are in **Case 1** of the preceeding section, then we need the induction basis for $X \setminus A$, which has only four elements, and we are done with that. Assume that it is **Case 2**. Denote the point in $X \setminus \{A, B, C, D\}$ by P. The angle $\angle PAD$ is smaller than α, since otherwise the paths $PADBC$ and $PADCB$ are good and thus at least one of them is R-avoiding. Analogously, the angle $\angle PBC$ is smaller than α. It follows that the path $CAPBD$ is good. Thus, its end directions $\overline{AC}$, $\overline{BD}$ are in the same R_i, say in R_1.

The paths $PABCD$ and $PBADC$ are good, so each of the opposite directions $\overline{CD}$, $\overline{DC}$ is in some R_i, say $\overline{CD} \in R_1$ and $\overline{DC} \in R_2$. Now R_1 contains $\overline{AC}$ and $\overline{CD}$, thus it contains also $\overline{AD}$. So R_1 contains $\overline{AD}$ and $\overline{BD}$. It follows that the angle $\angle ADB$ is at most $10\alpha < \pi/3$, contradicting that AB is a diameter of X.

Finally, if $|X| = 5$ is neither in **Case 1** nor in **Case 2**, then we use the induction basis on $X \setminus \{A, B, F\}$ which has two elements.

Acknowledgements. We thank Adrian Dumitrescu [2] and Géza Tóth for bringing this problem to our attention and for stimulating discussions. We thank Ferran Hurtado for pointing out reference [4] to us. Partial support for the first and second author from Hungarian National Foundation Grants No. 060427 and 062321 is also acknowledged. This work was also supported by the Discrete and Convex Geometry project, MTKD-CT-2005-014333, of the European Community. Research by the third author was also partially supported by project 1M0545 of the Ministry of Education of the Czech Republic.

References

[1] I. Bárány, A. Pór and P. Valtr. Paths with no small angles. In: E.S. Laber et al. (eds.), Theoretical Informatics. 8th Latin American Symposium, Búzios, Brasil, Springer, Lecture Notes in Computer Science 4957, 654–663 (2008).

[2] A. Dumitrescu. Personal communication, 2005.

[3] S. Fekete. Geometry and the Traveling Salesman Problem. PhD thesis, the University of Waterloo, 1992.

[4] S. Fekete and G.J. Woeginger, Angle-restricted tours in the plane, *Comput. Geom.: Theory and Appl.* **8** (1997), 195–218.

[5] D. Kirkpatrick and R. Seidel, The ultimate planar convex hull algorithm, *SIAM J. Computing*, **15** (1986), 286–297.