On the maximum size of an anti-chain of linearly separable sets and convex pseudo-discs[1]

Rom Pinchasi* Günter Rote†

January 11, 2008

Abstract

We answer a question raised by Walter Morris, and independently by Alon Efrat, about the maximum cardinality of an anti-chain composed of intersections of a given set of n points in the plane with half-planes. We approach this problem by establishing the equivalence with the problem of the maximum monotone path in an arrangement of n lines. A related problem on convex pseudo-discs is also discussed in the paper.

1 Introduction

Let P be a set of n points in the plane, no three of which are collinear. A subset of P is called *linearly separable* if it is the intersection of P with a closed half-plane. A k-set of P is a subset of k points from P which is linearly separable. Let $\mathcal{A}_k = \mathcal{A}_k(P)$ denote the collection of all k-sets of P. It is a well-known open problem to determine $f(k)$, the maximum possible cardinality of $\mathcal{A}_k$, where P varies over all possible sets of n points in general position in the plane. The current records are $f(k) = O(nk^{1/3})$ by Dey ([D98]) and $f(\lfloor n/2 \rfloor) \geq ne^{\Omega(\sqrt{\log n})}$ by Tóth ([T01]).

Let $\mathcal{A} = \mathcal{A}(P) = \cup_{k=0}^{n} \mathcal{A}_k$ be the family of all linearly separable subsets of P. The family $\mathcal{A}$ is partially ordered by inclusion. Clearly, each $\mathcal{A}_k$ is an anti-chain in $\mathcal{A}$. The following problem was raised by Walter Morris in 2003 in relation with the *convex dimension* of a point set (see [ES88]) and, as it turns out, it was independently raised by Alon Efrat 10 years before, in 1993:

Problem 1. What is the maximum possible cardinality $g(n)$ of an anti-chain in the poset $\mathcal{A}$, over all sets P with n points?

In Section 2 we show that in fact $g(n)$ can be very large, and in particular much larger than $f(n)$.

[1]This research was supported by a Grant from the G.I.F., the German-Israeli Foundation for Scientific Research and Development.

*Mathematics Dept., Technion—Israel Institute of Technology, Haifa 32000, Israel. `room@math.technion.ac.il`.

†Institut für Informatik, Freie Universität Berlin, Takustr. 9, 14195 Berlin, Germany. `rote@inf.fu-berlin.de`

Theorem 1. *$g(n) = \Omega(n^{2-\frac{d}{\sqrt{\log n}}})$, for some absolute constant $d > 0$.*

In an attempt to bound from above the function $g(n)$ one can view linearly separable sets as a special case of a slightly more general concept:

Definition 1. Let P be a set of n points in general position in the plane. A Family F of subsets of P is called a family of *convex pseudo-discs* if the following two conditions are satisfied:

1. Every set in F is the intersection of P with a convex set.
2. If A and B are two different sets in F, then both sets $\mathrm{conv}(A) \setminus \mathrm{conv}(B)$ and $\mathrm{conv}(B) \setminus \mathrm{conv}(A)$ are connected (or empty).

One natural example for a family of convex pseudo-discs is the family $\mathcal{A}(P)$, where P is a set of n points in general position in the plane. To see this, observe that every linearly separable set is the intersection of P with a convex set, namely, a half-plane. It is therefore left to verify that if $A = P \cap H_A$ and $B = P \cap H_B$, where H_A and H_B are two half-planes, then both $\mathrm{conv}(A) \setminus \mathrm{conv}(B)$ and $\mathrm{conv}(B) \setminus \mathrm{conv}(A)$ are connected. Let $A' = A \setminus H_B = A \setminus B = A \setminus \mathrm{conv}(B)$. Since $\mathrm{conv}(A') \cap \mathrm{conv}(B) = \emptyset$, we have $\mathrm{conv}(A) \setminus \mathrm{conv}(B) \supset \mathrm{conv}(A')$. For any $x \in \mathrm{conv}(A) \setminus \mathrm{conv}(B)$, we claim that there is a point $a' \in A'$ such that the line segment $[x, a']$ is fully contained in $\mathrm{conv}(A) \setminus \mathrm{conv}(B)$. This will clearly show that $\mathrm{conv}(A) \setminus \mathrm{conv}(B)$ is connected. Let a_1, a_2, a_3 be three points in A such that x is contained in the triangle $a_1a_2a_3$. If each line segment $[x, a_i]$, for $i = 1, 2, 3$, contains a point of $\mathrm{conv}(B)$, it follows that $x \in \mathrm{conv}(B)$, contrary to our assumption. Thus there must be a line segment $[x, a_i]$ that is contained in $\mathrm{conv}(A) \setminus \mathrm{conv}(B)$, and we are done.

In Section 3 we bound from above the maximum size of a family of convex pseudo-discs of a set P of n points in the plane, assuming that this family of subsets of P is by itself an anti-chain with respect to inclusion:

Theorem 2. *Let F be a family of convex pseudo-discs of a set P of n points in general position in the plane. If no member of F is contained in another, then F consists of at most $4\binom{n}{2} + 1$ members.*

Clearly, in view of Theorem 1, the result in Theorem 2 is nearly best possible. We show by a simple construction that Theorem 2 is in fact tight, apart from the constant multiplicative factor of n^2.

2 Large anti-chains of linearly separable sets

Instead of considering Problem 1 directly, we will consider a related problem.

Definition 2. For a pair x, y of points and a pair ℓ_1, ℓ_2 of non-vertical lines, we say that x, y *strongly separate* ℓ_1, ℓ_2 if x lies strictly above ℓ_1 and strictly below ℓ_2, and y lies strictly above ℓ_2 and strictly below ℓ_1.

We will also take the dual viewpoint and say that ℓ_1, ℓ_2 strongly separate x, y. (In fact, this relation is invariant under the standard point-line duality.)

If we have a set L of lines, we say that the point pair x, y is *strongly separated* by L, if L contains two lines ℓ_1, ℓ_2 that strongly separate x, y.

A pair of lines ℓ_1, ℓ_2 is said to be strongly separated by a set P of points if there are two points $x, y \in P$ that strongly separate ℓ_1 and ℓ_2.

Using the above terminology one can reduce Problem 1 to the following problem:

Problem 2. Let P be a set of n points in the plane. What is the maximum possible cardinality $h(n)$ (taken over all possible sets P of n points) of a set of lines L in the plane such that for every two lines $\ell_1, \ell_2 \in L$, P strongly separates ℓ_1 and ℓ_2.

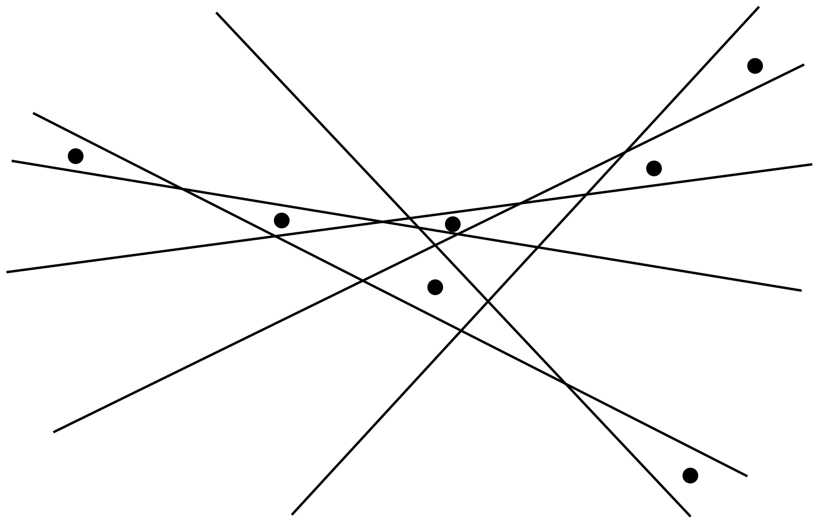

Figure 1: Problem 2.

To see the equivalence of Problem 1 and Problem 2, let P be a set of n points and L be a set of $h(n)$ lines that answer Problem 2. We can assume that none of the points lie on a line of L. Then with each of the lines $\ell \in L$ we associate the subset of P which is the intersection of P with the half-plane below ℓ. We thus obtain $h(n)$ subsets of P each of which is a linearly separable subset of P. Because of the condition on L and P, none of these linearly separable sets may contain another. Therefore we obtain $h(n)$ elements from $\mathcal{A}(P)$ that form an anti-chain, hence $g(n) \geq h(n)$.

Conversely, assume we have an anti-chain of size $g(n)$ in $\mathcal{A}(P)$ for a set P of n points. Each linearly separable set is the intersection of P with a half-plane, which is bounded by some line ℓ. We can assume without loss of generality that none of these lines is vertical, and at least half of the half-spaces lie below their bounding lines. These lines form a set L of at least $\lceil g(n)/2 \rceil$ lines, and each pair of lines is separated by two points from the n-point set P. Thus, $h(n) \geq \lceil g(n)/2 \rceil$.

Before reducing Problem 2 to another problem, we need the following simple lemma.

Lemma 1. *Let $\ell_1, \ldots, \ell_n$ be n non-vertical lines arranged in increasing order of slopes. Let P be a set of points. Assume that for every $1 \leq i < n$, P strongly separates ℓ_i and ℓ_{i+1}. Then for every $1 \leq i < j \leq n$, P strongly separates ℓ_i and ℓ_j.*

Proof. We prove the lemma by induction on $j - i$. For $j = i + 1$ there is nothing to prove. Assume $j - i \geq 2$. We first show the existence of a point $x \in P$ that lies above ℓ_i and below ℓ_j. Let B denote the intersection point of ℓ_i and ℓ_j. Let r_i denote the ray whose apex is B, included in ℓ_i, and points to the right. Similarly, let r_j denote the ray whose apex is B, included in ℓ_j, and points to the right.

Since the slope of ℓ_{i+1} is between the slope of ℓ_i and the slope of ℓ_j, ℓ_{i+1} must intersect either r_i or r_j (or both, in case it goes through B).

Case 1. ℓ_{i+1} intersects r_i. Then there is a point $x \in P$ that lies above ℓ_i and below ℓ_{i+1}. This point x must also lie below ℓ_j.

Case 2. ℓ_{i+1} intersects r_j. Then, by the induction hypothesis, there is a point $x \in P$ that lies above ℓ_{i+1} and below ℓ_j. This point x must also lie above ℓ_i.

The existence of a point y that lies above ℓ_j and below ℓ_i is symmetric. $\square$

By Lemma 1, Problem 2 is equivalent to following problem.

Problem 3. What is the maximum cardinality $h(n)$ of a collection of lines $L = \{\ell_1, \ldots, \ell_{h(n)}\}$ in the plane, indexed so that the slope of ℓ_i is smaller than the slope of ℓ_j whenever $i < j$, such that there exists a set P of n points that strongly separates ℓ_i and ℓ_{i+1}, for every $1 \leq i < h(n)$?

We will consider the dual problem of Problem 3:

Problem 4. What is the maximum cardinality $h(n)$ of a set of points $P = \{p_1, \ldots, p_{h(n)}\}$ in the plane, indexed so that the x-coordinate of p_i is smaller than the x-coordinate of p_j, whenever $i < j$, such that there exists a set L of n lines that strongly separates p_{i+1} and p_i, for every $1 \leq i < h(n)$?

We will relate Problem 4 to another well-known problem: the question of the longest monotone path in an arrangement of lines.

Consider an x-monotone path in a line arrangement in the plane. The *length* of such a path is the number of different line segments that constitute the path, assuming that consecutive line segments on the path belong to different lines in the arrangement. (In other words, if the path passes through a vertex of the arrangement without making a turn, this does not count as a new edge.)

Problem 5. What is the maximum possible length $\lambda(n)$ of an x-monotone path in an arrangement of n lines?

A construction of [BRSSS04] gives a simple line arrangement in the plane which consists of n lines and which contains an x-monotone path of length $\Omega(n^{2-\frac{d}{\sqrt{\log n}}})$ for some absolute constant $d > 0$. No upper bound that is asymptotically better than the trivial bound of $O(n^2)$ is known.

Problem 5 is closely related to Problem 4, and hence also to the other problems:

Proposition 1.

$$h(n) \geq \left\lceil \frac{\lambda(n) + 1}{2} \right\rceil, \tag{1}$$

$$\lambda(n) \geq h(n) - 2 \tag{2}$$

Proof. We first prove $h(n) \geq \lceil(\lambda(n)+1)/2\rceil$. Let L be a simple arrangement of n lines that admits an x-monotone path of length $m = \lambda(n)$. Denote by $x_0, x_1, \ldots, x_m$ the vertices of a monotone path arranged in increasing order of x-coordinates. In this notation $x_1, \ldots, x_{m-1}$ are vertices of the line arrangement L, while x_0 and x_m are chosen arbitrarily on the corresponding two rays which constitute the first and last edges, respectively, of the path. For each $1 \leq i < m$ let s_i denote the line that contains the segment $x_{i-1}x_i$, and let r_i denote the line through the segment $x_i x_{i+1}$.

For $1 \leq i < m$, we say that the path bends downward at the vertex x_i if the slope of s_i is greater than the slope of r_i, and it bends upward if the slope of s_i is smaller than the slope of r_i. Without loss of generality we may assume that at least half of the vertices $x_1, \ldots, x_{m-1}$ of the monotone path are downward bends.

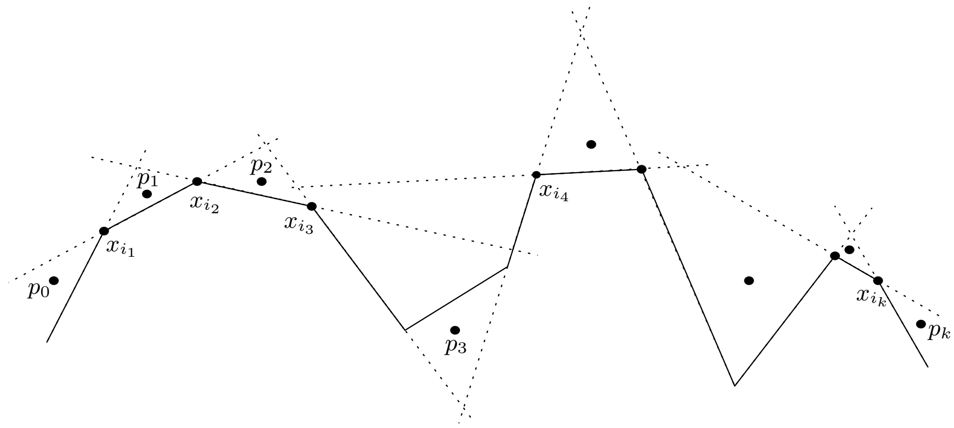

Figure 2: Constructing a solution for Problem 4.

Let $i_1 < i_2 < \cdots < i_k$ be all indices such that x_{i_j} is a downward bend, where $k \geq (m-1)/2$. Observe that for every $1 \leq j < k$, the monotone path between x_{i_j} and $x_{i_{j+1}}$ is an upward-bending convex polygonal path.

We will now define $k+1$ points $p_0, p_1, \ldots, p_k$ such that for every $0 \leq j < k$ the x-coordinate of p_j is smaller than the x-coordinate of p_{j+1}, and the line r_{i_j} lies above p_j and below p_{j+1} while the line s_{i_j} lies below p_j and above p_{j+1}. This construction will thus show that $h(n) \geq \lceil \frac{\lambda(n)+1}{2} \rceil$.

For every $1 \leq j \leq k$ let U_j and W_j denote the left and respectively the right wedges delimited by r_{i_j} and s_{i_j}. That is, U_j is the set of all points that lie below r_{i_j} and above s_{i_j}. Similarly, W_j is the set of all points that lie above r_{i_j} and below s_{i_j}.

Claim 1. *For every $1 \leq j < k$, W_j and U_{j+1} have a nonempty intersection.*

Proof. We consider two possible cases:

Case 1. $i_{j+1} = i_j + 1$. In this case $r_{i_j} = s_{i_{j+1}}$. Therefore any point above the line segment $[x_{i_j} x_{i_{j+1}}]$ that is close enough to that segment lies both below s_{i_j} and below $r_{i_{j+1}}$ and hence $W_j \cap U_{j+1} \neq \emptyset$.

Case 2. $i_{j+1} - i_j > 1$. In this case, as we observed earlier, the monotone path between x_{i_j} and $x_{i_{j+1}}$ is a convex polygonal path. Therefore, r_{i_j} and $s_{i_{j+1}}$ are different lines that meet at a point B whose x-coordinate is between the x-coordinates of x_{i_j} and $x_{i_{j+1}}$. Any point that lies vertically above B and close enough to B belongs to both W_j and U_{j+1}. $\square$

Now it is very easy to construct $p_0, p_1, \ldots, p_k$, see Figure 2. Simply take p_0 to be any point in U_1, and for every $1 \leq j < k$ let p_j be any point in $W_j \cap U_{j+1}$. Finally, let p_k be any point in W_k. It follows from the definition of $U_1, \ldots, U_k$ and $W_1, \ldots, W_k$ that for every $0 \leq j < k$, $r_{i_{j+1}}$ lies above p_j and below p_{j+1} and the line $s_{i_{j+1}}$ lies below p_j and above p_{j+1}.

We now prove the opposite direction: $\lambda(n) \geq h(n) - 2$.

Assume we are given $h(n)$ points $p_1, \ldots, p_{h(n)}$ sorted by x-coordinate and a set of n lines L such that every pair p_i, p_{i+1} is strongly separated by L. By perturbing the lines if necessary, we can assume that none of the lines goes through a point, and no three lines are concurrent. For $1 < i < h(n)$, let f_i be the face of the arrangement that contains p_i, and let A_i and B_i be, respectively, the left-most and right-most vertex in this face. (The faces f_i are bounded, and therefore A_i and B_i are well-defined.) The monotone path will follow the upper boundary of each face f_i from A_i to B_i.

We have to show that we can connect B_i to A_{i+1} by a monotone path. This follows from the separation property of L. Let s_i, r_i be a pair of lines that strongly separates p_i and p_{i+1} in such a way that r_i lies above p_i and below p_{i+1} and s_i lies below p_i and above p_{i+1}. Since B_i lies on the boundary of the face f_i that contains p_i, B_i lies also between r_i and s_i, including the possibility of lying on these lines. We can thus walk on the arrangement from B_i to the right until we hit r_i or s_i, and from there we proceed straight to the intersection point Q_i of r_i and s_i. Similarly, there is a path in the arrangement from A_{i+1} to the left that reaches Q_i. and these two paths together link B_i with A_{i+1}.

To count the number of edges of this path, we claim that there must be at least one bend between B_i and A_{i+1} (including the boundary points B_i and A_{i+1}). If there is no bend at Q_i, the path must go straight through Q_i, say, on r_i. But then the path must leave r_i at some point when going to the right: if the path has not left r_i by the time it reaches A_{i+1} and A_{i+1} lies on r_i, then the path must bend upward at this point, since it proceeds on the upper boundary of the face f_{i+1} that lies above r_i.

Thus, the path makes at least $h(n)-3$ bends (between B_i and A_{i+1}, for $1 < i < h(n)-1$) and contains at least $h(n) - 2$ edges. $\square$

Now it is very easy to give a lower bound for $g(n)$, and prove Theorem 1. Indeed, this follows because $g(n) \geq h(n)$ and $h(n) \geq \lceil \frac{\lambda(n)+1}{2} \rceil = \Omega(n^{2-\frac{d}{\sqrt{\log n}}})$,

The close relation between Problems 1 and 5 comes probably as no big surprise if one considers the close connection between k-sets and *levels* in arrangements of lines (see [E87, Section 3.2]). For a given set of n points P, the k-sets are in one-to-one correspondence with the faces of the dual arrangements of lines which have k lines passing below them and $n-k$ lines passing above them (or vice versa). The lower boundaries of these cells form the k-th level in the arrangement, and the upper boundaries form the $(k+1)$-st level.

Our chain of equivalence from Problem 1 to Problem 5 extends this relation between k-sets and levels in a way that is not entirely trivial: for example, establishing that we get sets that form an antichain requires some work, whereas for k-sets this property is fulfilled automatically.

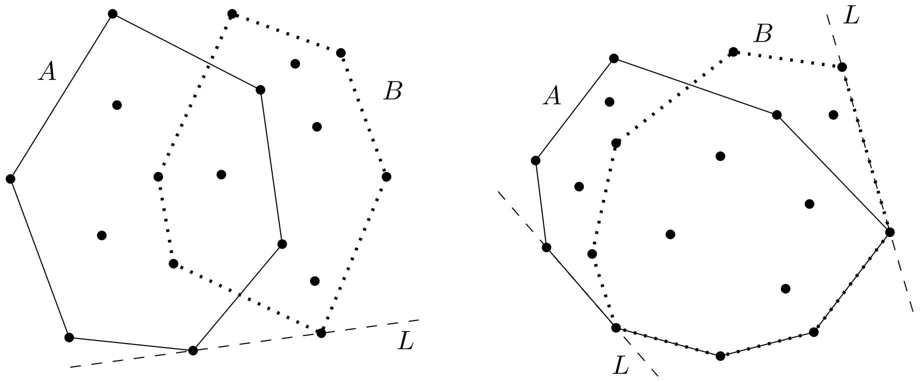

Figure 3: The two cases of common tangents in Lemma 2

3 Proof of Theorem 2

The heart of our argument uses a linear algebra approach first applied by Tverberg [T82] in his elegant proof for a theorem of Graham and Pollak [GP72] on decomposition of the complete graph into bipartite graphs.

Let F be a collection of convex pseudo-discs of a set P of n points in general position in the plane. We wish to bound from above the size of F assuming that no set in F contains another. For every directed line $L = \overrightarrow{xy}$ passing through two points x and y in P we denote by L_x the collection of all sets $A \in F$ that lie in the closed half-plane to the left of L such that L touches $\text{conv}(A)$ at the point x only. Similarly, let L_y be the collection of all sets $A \in F$ that lie in the closed half-plane to the left of L such that L touches $\text{conv}(A)$ at the point y only. Finally, let L_{xy} be those sets $A \in F$ that lie in the closed half-plane to the left of L such that L supports $\text{conv}(A)$ at the edge xy.

Definition 3. Let A and B be two sets in F. Let L be a directed line through two points x and y in P. We say that L is a common tangent of the *first kind* with respect the pair (A, B) if $A \in L_x$ and $B \in L_y$.

We say that L is a common tangent of the second kind with respect to (A, B) if $A \in L_{xy}$ and $B \in L_y$, or if $A \in L_x$ and $B \in L_{xy}$.

The crucial observation about any two sets A and B in F is stated in the following lemma.

Lemma 2. *Let A and B be two sets in F. Then exactly one of the following conditions is true.*

1. *There is precisely one common tangent of the first kind with respect to (A, B) and no common tangent of the second kind with respect to (A, B), or*

2. *there is no common tangent of the first kind with respect to (A, B), and there are precisely two common tangents of the second kind with respect (A, B).*

Proof. The idea is that because A and B are two pseudo-discs and none of $\text{conv}(A)$ and $\text{conv}(B)$ contains the other, then as we roll a tangent around $C = \text{conv}(A \cup B)$, there is precisely one transition between A and B, and this is where the situation described in the lemma occurs (see Figure 3).

Formally, by our assumption on F, none of A and B contains the other. Any directed line L that is a common tangent of the first or second kind with respect to A and B must be a line supporting $\mathrm{conv}(A \cup B)$ at an edge.

Let $x_0, \ldots, x_{k-1}$ denote the vertices of $C = \mathrm{conv}(A \cup B)$ arranged in counterclockwise order on the boundary of C. In what follows, arithmetic on indices is done modulo k.

There must be an index i such that $x_i \in A \setminus B$, for otherwise every x_i belongs to B and therefore $\mathrm{conv}(B) = \mathrm{conv}(A \cup B) \supset \mathrm{conv}(A)$ and therefore $B \supset A$ (because both A and B are intersections of P with convex sets) in contrast to our assumption. Similarly, there must be an index i such that $x_i \in B \setminus A$.

Let I_A be the set of all indices i such that $x_i \in A \setminus B$, and let I_B be the set of all indices i such that $x_i \in B \setminus A$.

We claim that I_A (and similarly I_B) is a set of consecutive indices. To see this, assume to the contrary that there are indices i, j, i', j' arranged in a cyclic order modulo k such that $x_i, x_{i'} \in A \setminus B$ and $x_j, x_{j'} \in B$. Then it is easy to see that $\mathrm{conv}(A) \setminus \mathrm{conv}(B)$ is not a connected set because x_i and $x_{i'}$ are in different connected components of this set.

We have therefore two disjoint intervals $I_A = \{i_A, i_A + 1, \ldots, j_A\}$ and $I_B = \{i_B, i_B + 1, \ldots, j_B\}$. It is possible that $i_A = j_A$ or $i_B = j_B$.

Observe that $x_{i_A}, x_{j_A}, x_{i_B}, x_{j_B}$ are arranged in this counterclockwise cyclic order on the boundary of C, and for every index $i \notin I_A \cup I_B$, $x_i \in A \cap B$. The only candidates for common tangents of the first kind or of the second kind with respect to A and B are of the form $\overrightarrow{x_i x_{i+1}}$, that is, they must pass through two consecutive vertices of C.

We distinguish two possible cases:

1. $i_B = j_A + 1$. In this case the line through x_{j_A} and x_{i_B} is the only common tangent of the first kind with respect to (A, B) and there are no common tangents of the second kind with respect to (A, B).

2. $i_B \neq j_A + 1$. In this case, there is no common tangent of the first kind with respect to (A, B). The line through $x_{i_B - 1}$ and x_{i_B} and the line through x_{j_A} and x_{j_A+1} are the only common tangents of the second kind with respect to (A, B).

This completes the proof of the lemma. $\square$

Let $A_1, \ldots, A_m$ be all the sets in F, and for every $1 \leq i \leq m$ let z_i be an indeterminate associated with A_i. For each directed line $L = \overrightarrow{xy}$, define the following polynomial P_L:

$$P_L(z_1, \ldots, z_m) = \Big(\sum_{A_i \in L_x} z_i\Big)\Big(\sum_{A_j \in L_y} z_j\Big) + \frac{1}{2}\Big(\sum_{A_i \in L_x} z_i\Big)\Big(\sum_{A_j \in L_{xy}} z_j\Big) + \frac{1}{2}\Big(\sum_{A_i \in L_y} z_i\Big)\Big(\sum_{A_j \in L_{xy}} z_j\Big)$$

This polynomial contains a term $z_u z_v$ whenever L is a tangent line for the pair (A_u, A_v) or for the pair (A_v, A_u) (of the first or of the second kind, and with coefficient 1 or $\frac{1}{2}$, accordingly). If we sum this equation over all directed lines L, it follows by Lemma 2 that every term $z_u z_v$ with $u \neq v$ appears with coefficient 2:

$$\sum_L P_L(z_1, \ldots, z_m) = \sum_{u<v} 2 z_u z_v = (z_1 + \cdots + z_m)^2 - (z_1^2 + \cdots + z_m^2) \qquad (3)$$

Consider the system of linear equations $\sum_{A_i \in L_x} z_i = 0$ and $\sum_{A_i \in L_y} z_i = 0$, where $L = \overrightarrow{xy}$ varies over all directed lines determined by P. Add to this system the equation $z_1 + \cdots + z_m = 0$. There are $4\binom{n}{2} + 1$ equations in this system and if $m > 4\binom{n}{2} + 1$, there must be a nontrivial solution. However, it is easily seen that a nontrivial solution $(z_1, \ldots, z_m)$ will result in a contradiction to (3). This is because the left-hand side of (3) vanishes, while the right-hand side equals $-(z_1^2 + \cdots + z_m^2) \neq 0$. We conclude that $|F| = m \leq 4\binom{n}{2} + 1$. $\square$

We now show by a simple construction that Theorem 2 is tight apart from the multiplicative constant factor of n^2. Fix three rays r_1, r_2, and r_3 emanating from the origin such that the angle between two rays is 120 degrees. For each $i = 1, 2, 3$, let $p_1^i, \ldots, p_n^i$ be n points on r_i, indexed according to their increasing distance from the origin. Slightly perturb the points to get a set P of $3n$ points in general position in the plane. For every $1 \leq j, k, l \leq n$ define

$$F_{jkl} = \{p_1^1, \ldots, p_j^1\} \cup \{p_1^2, \ldots, p_k^2\} \cup \{p_1^3, \ldots, p_l^3\}.$$

It can easily be checked that the collection of all F_{jkl} such that $1 \leq j, k, l \leq n$ and $j + k + l = n + 2$ is an anti-chain of convex pseudo-discs of P. This collection consists of $\binom{n+1}{2}$ sets.

References

[BRSSS04] J. Balogh, O. Regev, C. Smyth, W. Steiger, and M. Szegedy, Long monotone paths in line arrangements. *Discrete Comput. Geom.* **32** (2004), no. 2, 167–176.

[D98] T. K. Dey, Improved bounds for planar k-sets and related problems. *Discrete Comput. Geom.* **19** (1998), no. 3, 373–382.

[ES88] P. H. Edelman and M. E. Saks, Combinatorial representation and convex dimension of convex geometries. *Order* **5** (1988), no. 1, 23–32.

[E87] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, EATCS Monographs on Theoret. Comput. Sci., vol. 10, Springer-Verlag, Berlin, 1987.

[GP72] R. L. Graham and H. O. Pollak, On embedding graphs in squashed cubes. In *Proc. Conf. Graph Theory Appl.*, Western Michigan Univ., May 10–13, 1972, ed. Y. Alavi, D. R. Lick, and A. T. White, Lecture Notes in Mathematics, vol. 303, Springer-Verlag, Berlin, 1972, pp. 99–110.

[T01] G. Tóth, Point sets with many k-sets, *Discrete Comput. Geom.* **26** (2001) no. 2, 187–194.

[T82] H. Tverberg, On the decomposition of K_n into complete bipartite graphs. *J. Graph Theory* **6** (1982), no. 4, 493–494.