

On k -intersecting curves and related problems

Sarit Buzaglo* Ron Holzman† Rom Pinchasi‡

December 4, 2007

Abstract

Let P be a set of n points in the plane and let \mathcal{C} be a family of simple closed curves in the plane each of which avoids the points of P . For every curve $C \in \mathcal{C}$ we denote by $\text{disc}(C)$ the region in the plane bounded by C . Fix an integer $k \geq 0$ and assume that every two curves in \mathcal{C} intersect at most k times and that for every two curves $C, C' \in \mathcal{C}$ the intersection $\text{disc}(C) \cap \text{disc}(C')$ is a connected set. We consider the family $\mathcal{F} = \{P \cap \text{disc}(C) \mid C \in \mathcal{C}\}$. When k is even, we provide sharp bounds, in terms of n, k , and ℓ , for the number of sets in \mathcal{F} of cardinality ℓ , assuming that $\bigcap_{C \in \mathcal{C}} \text{disc}(C)$ is nonempty. In particular, we provide sharp bounds for the number of halving pseudo-parabolas for a set of n points in the plane. Finally, we consider the VC-dimension of \mathcal{F} and show that \mathcal{F} has VC-dimension at most $k + 1$.

1 Introduction

Let C be a simple closed Jordan curve in the plane. By Jordan's Theorem C divides the plane into two regions, only one of which is bounded. We call the bounded region the *disc* bounded by C and we denote this region by $\text{disc}(C)$. Any point p inside $\text{disc}(C)$ is said to be *surrounded* by C and C is said to be *surrounding* p .

A *bi-infinite x -monotone* curve is a curve that crosses every vertical line at precisely one point. Any graph of a continuous function defined on the real numbers is an example for such a curve. If \mathcal{C} is a family of bi-infinite x -monotone curves, every two of which cross an even number of times, then it is easy to see that \mathcal{C} can be realized as a family of simple closed curves with the connected-intersection property. This can be done by identifying the two ends at infinity for each curve in \mathcal{C} .

Any arrangement of lines in the plane is an example for a 1-intersecting family of bi-infinite x -monotone curves. In fact, any arrangement of x -monotone pseudo-lines is yet another such example.

Given a set P of n points in general position in the plane. An ℓ -set of P is a subset of ℓ points from P which is the intersection of P with a closed half-plane. It is a well-known open problem to determine $f(\ell, n)$, the maximum possible cardinality of a family of ℓ -sets of a set P of n points in the plane. The current records are $f(\ell, n) = O(n\ell^{1/3})$ by Dey ([4]) and $f(k, n) \geq ne^{\Omega(\sqrt{\log k})}$ by Tóth ([10]).

*Mathematics Dept., Technion—Israel Institute of Technology, Haifa 32000, Israel.
sarahb@tx.technion.ac.il

†Mathematics Dept., Technion—Israel Institute of Technology, Haifa 32000, Israel.
holzman@tx.technion.ac.il

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

This notion of an ℓ -set can be easily generalized for any collection of bi-infinite x -monotone curves with respect to a set P of n points in the plane. Let P be a set of n points in the plane and let \mathcal{C} be a family of bi-infinite x -monotone curves. Call a subset $S \subset P$ of cardinality ℓ an ℓ -set of P with respect to \mathcal{C} , if there is a curve in \mathcal{C} that lies above each point of S and below each point of $P \setminus S$.

In fact, Dey's bound of $O(n\ell^{1/3})$ is a valid bound for the number of ℓ -sets of a set of n points with respect to any family \mathcal{C} of 1-intersecting bi-infinite x -monotone curves (that is, x -monotone pseudo-lines). Surprisingly, we can provide sharp bounds to the number of ℓ -sets of a set P of n points in the plane with respect to any family \mathcal{C} of k -intersecting bi-infinite x -monotone curves, for k even (Theorem 1). As we shall remark after the proof of Theorem 1, using Dey's bound and the (inductive) proof of Theorem 1, we can provide 'good' bounds for the number of ℓ -sets above also when k is odd.

Given a finite set of points P in the plane and a simple closed curve C in the plane that avoids the points of P , we denote by P_C the set $P \cap \text{disc}(C)$.

Definition 1.1. *Let \mathcal{C} be a family of simple closed curves in the plane. We say that \mathcal{C} has the k -intersection property if any two curves in \mathcal{C} intersect properly in at most k points. We say that \mathcal{C} has the connected intersection property if for every $C, C' \in \mathcal{C}$ the set $\text{disc}(C) \cap \text{disc}(C')$ is either connected or empty.*

Let \mathcal{F} be a family of subsets of $\{1, \dots, n\}$. A subset $S \subset \{1, \dots, n\}$ is said to be *shattered* by \mathcal{F} if for any subset B of S there exists $F \in \mathcal{F}$ with $B = F \cap S$. The *VC-dimension* ([11]) of \mathcal{F} is the largest cardinality of a set S that \mathcal{F} *shatters*.

Perhaps one of the most fundamental results on VC-dimension is the Perles-Sauer-Shelah theorem ([7, 8]), which says that a family \mathcal{F} of subsets of $\{1, \dots, n\}$ that has VC-dimension d consists of at most $\binom{n}{0} + \dots + \binom{n}{d} = O(n^d)$ members.

In Section 3 we study the VC-dimension of the family $\mathcal{F} = \{P_C \mid C \in \mathcal{C}\}$ for a fixed set of n points P in the plane and a family \mathcal{C} of simple closed curves which has both the k -intersection property, for some fixed $k \geq 2$, and the connected intersection property. We show that the VC-dimension of such a family is at most $k + 1$ (Theorem 7).

2 k -intersecting bi-infinite x -monotone curves

In this section we prove the following main theorem:

Theorem 1. *Let $k \geq 0$ be a fixed even number. Let P be a set of n points in the plane and let \mathcal{C} be a family of bi-infinite x -monotone curves, every two of which intersect at most k times. Assume that no curve in \mathcal{C} passes through a point of P . Then for every $\ell \leq \frac{n}{2}$, P has at most $O(\ell^{k/2} n^{k/2})$ ℓ -sets with respect to \mathcal{C} . This bound is best possible.*

In order to prove Theorem 1, we perform a small perturbation of the points of P so that no two points of P have the same x -coordinate. We assign to each curve in \mathcal{C} a vector in $\{0, 1\}^n$ as follows. Let p_1, \dots, p_n denote the points of P ordered according to the increasing order of their x -coordinates. To a curve $C \in \mathcal{C}$ we assign the vector v_C whose i 'th coordinate is 0 if p_i lies above C , and is equal to 1 if p_i lies below C .

For a set of vectors $V \subset \{0, 1\}^m$, we say that V has the t -intersection property if there are no two vectors $u_1, u_2 \in V$ and $t + 2$ indices $1 \leq j_1 < \dots < j_{t+2} \leq m$ such that for every $1 \leq s \leq t + 2$ the j_s 'th coordinate of u_1 equals 1 if s is odd and equals 0 if s is even, and the j_s 'th coordinate of u_2 equals 0 if s is odd and equals 1 if s is even.

Observe that because the family of curves \mathcal{C} has the k -intersection property, then the set of vectors $\{v_C \mid C \in \mathcal{C}\}$ must have the k -intersection property.

Theorem 1 will therefore follow from the following theorem on sets of vectors in $\{0, 1\}^n$.

Theorem 2. *Let $V \subset \{0, 1\}^n$ be a set of vectors which has the k -intersection property for some fixed even number $k \geq 0$. Assume further that every vectors in V has precisely ℓ 1-entries, where $\ell \leq \frac{n}{2}$. Then $|V| = O(\ell^{k/2} n^{k/2})$. This bound is best possible.*

Proof. We prove the theorem by induction on k . If $k = 0$, then clearly V may consist of at most one vector. Assume that $k \geq 2$.

For a vector $v \in V$ denote by v^i the vector in $\{0, 1\}^i$ which consists of the first i coordinates of v . For every $1 \leq i \leq n$ denote by T_i the set $T_i = \{v^i \mid v \in V\}$. Observe that if u is a vector in T_{i+1} , then the vector in $\{0, 1\}^i$ obtained from u by removing its $(i+1)$ 'th coordinate is a member of T_i . This is because if $u \in T_{i+1}$, then $u = v^{i+1}$ for some $v \in V$. Hence $u^i = v^i \in T_i$.

Since we are interested in bounding the cardinality of $T_n = V$, let us consider how large can $|T_{i+1}|$ be compared to $|T_i|$. Define a bipartite graph H whose vertices are the vectors in T_i and the vectors in T_{i+1} . We connect a vector $u \in T_i$ to a vector $v \in T_{i+1}$ if the first i coordinates of v are identical with those of u , that is, if $v^i = u$. Clearly, every vector $u \in T_i$ is connected in H to either one or two vectors in T_{i+1} . Denote by T'_i the set of vectors in T_i that are connected in H to two vectors in T_{i+1} . Observe that we always have $|T_{i+1}| = |T_i| + |T'_i|$. This is because every vector $v \in T_{i+1}$ is connected in H to precisely one vector, namely v^i , in T_i . We claim that $|T'_i| = O(\ell^{\frac{k}{2}} n^{\frac{k}{2}-1})$ for every $1 \leq i \leq n-1$. To see this we show that for every fixed $0 \leq j \leq \ell$ the subset of T'_i which consists of all those vectors having precisely j 1-entries, consists of $O(\ell^{\frac{k}{2}-1} n^{\frac{k}{2}-1})$ vectors. Since the number of 1-entries in each vector in T'_i is at most ℓ , the assertion $|T'_i| = O(\ell^{\frac{k}{2}} n^{\frac{k}{2}-1})$ will follow.

Therefore, fix j between 1 and ℓ and let B be the set of all vectors in T'_i with precisely j 1-entries. We claim that B has the $(k-2)$ -intersection property. Once we establish this, it will follow from the induction hypothesis that $|B| = O(\ell^{\frac{k}{2}-1} n^{\frac{k}{2}-1})$ as required.

Assume to the contrary that B does not have the $(k-2)$ -intersection property. This means that there are two vectors u_1 and u_2 in $B \subset T'_i$ and k indices $1 \leq j_1 < \dots < j_k \leq i$ such that for every $1 \leq s \leq k$ the j_s 'th coordinate of u_1 equals 1 if s is odd and equals 0 if s is even, and the j_s 'th coordinate of u_2 equals 0 if s is odd and equals 1 if s is even.

As $u_1, u_2 \in T'_i$, both vectors $(u_1, 0)$ and $(u_2, 1)$ are in T_{i+1} . Let $v_1 \in V$ be the vector such that $v_1^{i+1} = (u_1, 1)$ and let $v_2 \in V$ be the vector such that $v_2^{i+1} = (u_2, 0)$. Observe that there is one more 1-entry in v_1^{i+1} than in v_2^{i+1} . Therefore, as the number of 1-entries in both v_1 and v_2 is ℓ , it follows that there must be an index j_{k+1} such that $i+1 < j_{k+1} \leq n$ and the j_{k+1} 'th coordinate of v_1 is 0 while the j_{k+1} coordinate of v_2 is 1. Considering the coordinates $j_1, \dots, j_k, i+1, j_{k+1}$ of v_1 and v_2 , we see that V does not have the k -intersection property, which is a contradiction.

It now follows that

$$|T_{i+1}| = |T_i| + |T'_i| = |T_i| + O(\ell^{\frac{k}{2}} n^{\frac{k}{2}-1})$$

for every $1 \leq i < n$. Hence, $|V| = |T_n| = O(\ell^{\frac{k}{2}} n^{\frac{k}{2}})$, as desired.

To see that the above bound on V is best possible in terms of ℓ and n , let V be the set of all vectors $v \in \{0, 1\}^n$ with the following properties:

1. The first coordinate of v equals 1.

2. v has precisely ℓ coordinates that are equal to 1.
3. k is the maximum even number such that there exist $k+1$ indices $1 \leq j_1 < \dots < j_{k+1} \leq n$ such that for every odd s between 1 and $k+1$ the j_s 'th coordinate of v equals 1 and for every even s between 1 and $k+1$ the j_s 'th coordinate of v equals 0.

We show that $|V| = \Omega(\ell^{\frac{k}{2}} n^{\frac{k}{2}})$ and that V has the k -intersection property.

Note that the conditions on the vectors in V are equivalent to the condition that each vector in V is composed from $k/2 + 1$ blocks of consecutive entries that equal to 1. Every two such blocks are separated by at least one 0-entry. In addition the leftmost block includes the position of the first coordinate, and the total length of all blocks of 1-entries is ℓ .

It is now an elementary problem in enumerative combinatorics to determine the cardinality of V precisely. Indeed, there are precisely $\binom{\ell-1}{k/2}$ ways to decide about the lengths of the 1-entries blocks (that in total sum up to a length of ℓ). Then we just need to decide about how many 0-entries are there between any two blocks of 1-entries, keeping in mind that the total length of the vector is n . This can be done in exactly $\binom{n-\ell}{k/2}$ ways. Therefore the cardinality of V is exactly $\binom{\ell-1}{k/2} \binom{n-\ell}{k/2} = \Omega(\ell^{k/2} n^{k/2})$.

It is left to show that V has the k -intersection property. Assume not, then there are two vectors $v_1, v_2 \in V$ and $k+2$ indices $1 \leq j_1 < \dots < j_{k+2} \leq n$ such that for every $1 \leq s \leq k+2$ the j_s 'th coordinate of v_1 equals 1 if s is odd and equals 0 if s is even, and the j_s 'th coordinate of v_2 equals 0 if s is odd and equals 1 if s is even. We will concentrate on v_2 . Observe that j_1 is necessarily greater than 1 because, as $v_2 \in V$, we know that the first coordinate of v_2 is equal to 1. Now, set $j_0 = 1$ and consider the indices $j_0 < j_1 < \dots < j_{k+2}$. We know that the j_0, j_2, \dots, j_{k+2} entries in v_2 are all equal to 1 while the j_1, j_3, \dots, j_{k+1} entries in v_2 are all equal to 0. This means that v_2 violates condition (3) in the requirements on the vectors in V . We have thus reached the desired contradiction. ■

Remark 1. It is not difficult and in fact rather straight forward to show that the construction in Theorem 2 which shows that the bound $O(\ell^{k/2} n^{k/2})$ is best possible, can be used to yield a construction of a family \mathcal{C} of bi-infinite x -monotone curves with the k -intersection property and a set P of n points in the plane such that P has $\Omega(\ell^{k/2} n^{k/2})$ ℓ -sets with respect to \mathcal{C} . We will omit the details and just sketch the construction. Let P be the set of points $(1, 0), (2, 0), \dots, (n, 0)$ on the x -axis. Let V be a set of vectors in $\{0, 1\}^n$ of cardinality $\Omega(\ell^{\frac{k}{2}} n^{\frac{k}{2}})$ which has the k -intersection property such that each vector in V has precisely ℓ 1-entries. For every vector $v \in V$ we construct a bi-infinite x -monotone curve C such that $v = v_C$. This can be done by drawing a bi-infinite x -monotone curve that goes above the point $(i, 0)$, if the i 'th coordinate of v equals 1 and which goes below $(i, 0)$, if the i 'th coordinate of v equals 0. The fact that V has the k -intersection property can be used to show that the set of curves \mathcal{C} thus constructed has the k -intersection property if we avoid *unnecessary* crossings between the constructed curves.

Remark 2. In view of the lower bounds known for the maximum number of ℓ -sets with respect to a family of 1-intersecting bi-infinite x -monotone curve, that is, for the function $f(\ell, n)$ introduced earlier, it is evident that the bounds of Theorem 1 are not valid when k is odd. The crucial point where the proof of Theorem 1 will collapse is the basis of induction, namely, the case $k = 1$. However, using Dey's upper bound ([4]) for $f(\ell, n)$ we can apply the induction step in the proof of Theorem 1 in the case when k is odd and obtain the following result:

Theorem 3. *Let $k \geq 1$ be a fixed odd number. Let P be a set of n points in the plane and let \mathcal{C} be a family of bi-infinite x -monotone curves, every two of which intersect at most k times. Assume that no curve in \mathcal{C} passes through a point of P . Then P has at most*

$$f(\ell, n)O(\ell^{\frac{k-1}{2}}n^{\frac{k-1}{2}}) = O(\ell^{1/3}n\ell^{\frac{k-1}{2}}n^{\frac{k-1}{2}}) = O(\ell^{(\frac{k}{2}-\frac{1}{6})}n^{\frac{k+1}{2}})$$

ℓ -sets with respect to \mathcal{C} .

Returning to Theorem 1, we now can immediately draw some simple corollaries in the case $k = 2$.

A family \mathcal{C} of bi-infinite x -monotone curves is called a family of *pseudo-parabolas* if every two curves in \mathcal{C} are either disjoint, or properly cross in at most two points. In other words, a family of pseudo-parabolas is nothing else but a family of bi-infinite x -monotone curves with the 2-intersection property.

A family \mathcal{C} of simple closed curves in the plane is called a family of *pseudo-circles* if every two curves in \mathcal{C} are either disjoint, or properly cross at precisely two points.

By a result of Snoeyink and Hershberger ([9]), any family of pseudo-circles surrounding a common point can be *swept by a ray*. In other words, it can be realized as a family of 2-intersecting bi-infinite x -monotone curves (see [9] for the formal definition of a sweeping) and this can be done by a one to one continuous transformation of the plane, after we identify the two ends at infinity of each curve.

Hence, we immediately get the following theorem:

Theorem 4. *Let P be a set of n points in the plane. Let \mathcal{C} be a family of pseudo-circles that avoid the points of P . Assume that there is a point o which belongs to $\text{disc}(C)$ for every $C \in \mathcal{C}$ and that each $C \in \mathcal{C}$ surround precisely ℓ points of P . If no two curves in \mathcal{C} surround the same set of points of P , then $|\mathcal{C}| = O(\ell n)$.*

As a corollary we get the same bound under somewhat weaker conditions on \mathcal{C} , as follows.

Corollary 5. *Let P be a set of n points in the plane. Let \mathcal{C} be a family of pseudo-circles that avoid the points of P . Assume that each curve in \mathcal{C} surround precisely ℓ points of P and that every two curves in \mathcal{C} properly cross. If no two curves in \mathcal{C} surround the same set of points of P , then $|\mathcal{C}| = O(\ell n)$.*

Proof. We need the following easy observation proved in [1].

Lemma 6. *Among any five pseudo-discs bounded by the elements of \mathcal{C} , there are at least three that have a point in common.*

Let $p \geq q \geq 2$ be integers. We say that a family F of sets has the (p, q) *property* if among every p members of F there are q that have a point in common. We say that a family of sets F is *pierced* by a set T if every member of F contains at least one element of T . The set T is often called a *transversal* of F . Fix $p \geq q \geq d + 1$. Alon and Kleitman [3] proved that there exists a transversal of size at most $r = r(p, q, d)$ for any finite family of convex sets in \mathbb{R}^d with the (p, q) -property. In [2], this result was extended to any finite family F of open regions in d -space with the property that the intersection of every subfamily of F is either empty or contractible. Their result implies the following. There is an absolute constant r such that any family of discs bounded by pairwise intersecting pseudo-circles can be pierced by at most r points.

Now fix a set $\{o_1, o_2, \dots, o_r\}$ of r points that pierces $\text{disc}(C)$ for every $C \in \mathcal{C}$. Let \mathcal{C}_i consist of all elements of \mathcal{C} that surround o_i , for $i = 1, 2, \dots, r$. From Theorem 4 it follows that $|\mathcal{C}_i| = O(\ell n)$ for every $1 \leq i \leq r$. Hence $|\mathcal{C}| \leq |\mathcal{C}_1| + \dots + |\mathcal{C}_r| = O(\ell n)$. ■

3 VC-dimension of k -intersecting curves

In this section we prove the following theorem:

Theorem 7. *Let P be a set of n points in the plane and let \mathcal{C} be a family of simple closed curves avoiding the points of P . Assume that \mathcal{C} has both the k -intersection property, for some fixed $k \geq 2$, and the connected intersection property. Then the family $\mathcal{F} = \{P_C \mid C \in \mathcal{C}\}$ has VC-dimension at most $k + 1$.*

For the proof of Theorem 7 we will need some preliminary results. The next lemma is a generalization of Helly's theorem ([5]) founded by Molnár ([6]):

Lemma 8. *Any finite family of at least three regions in the plane has a nonempty simply connected intersection, provided that any two of its members have a connected intersection and any three have a nonempty intersection.*

We will need also the following lemma that can be found in [1]:

Lemma 9. *Let \mathcal{D} be a family of closed curves such that, any pair of discs bounded by curves in \mathcal{D} has either an empty intersection or a connected intersection. Assume that all the curves in \mathcal{D} have a common interior point O . Then the union of any set of discs bounded by curves in \mathcal{D} is simply connected.*

Before getting to the proof of Theorem 7, we need one more crucial lemma:

Lemma 10. *Let \mathcal{D} be a finite family of closed curves. Assume that the union of any number of discs bounded by curves in \mathcal{D} is simply connected. Let y be an arbitrary point in $\mathbb{R}^2 \setminus \cup_{C \in \mathcal{D}} C$. Consider the family $\mathcal{D}_y \subseteq \mathcal{D}$ of all the curves in \mathcal{D} which surround y . Then there exists a Jordan arc, connecting y to a point at infinity, which intersects every curve in \mathcal{D}_y exactly once and avoids all the curves in $\mathcal{D} \setminus \mathcal{D}_y$.*

Proof. We shall prove the lemma by induction on $|\mathcal{D}_y|$. The case $|\mathcal{D}_y| = 0$ is trivial. Suppose $|\mathcal{D}_y| > 0$. The induction hypothesis states that for any point $p \in \mathbb{R}^2 \setminus \cup_{C \in \mathcal{D}} C$ with $|\mathcal{D}_p| < |\mathcal{D}_y|$, there exists an arc, connecting p to a point at infinity, which intersects every curve in \mathcal{D}_p exactly once and avoids all the curves in $\mathcal{D} \setminus \mathcal{D}_p$. The arrangement of curves in \mathcal{D} can be viewed as a drawing of a planar graph with a vertex set V , consisting of all the intersection points of curves in \mathcal{D} , together with a set of edges E , consisting of all the connected components in $\cup_{C \in \mathcal{D}} C \setminus V$. There exists a face F_y of this arrangement which contains y . The face F_y must be bounded since $|\mathcal{D}_y| > 0$. An edge of F_y will be called an *inner edge* if it is a portion of a curve in \mathcal{D}_y . We claim that F_y must have an inner edge. To see this, assume to the contrary that F_y does not have an inner edge. Consider the set of all curves in \mathcal{D} which contain an edge of F_y and let U be the union of all the discs bounded by these curves. By our assumption, U is a simply connected region. Observe that $y \notin U$, and any arc from y to infinity must cross U . Thus $\mathbb{R}^2 \setminus U$ is not connected, hence U is not simply connected, which yields a contradiction. We conclude that F_y must have an inner edge.

Let us choose an inner edge of F_y and draw an arc γ , starting at y , which crosses the inner edge once and does not cross any other curve. Denote by x the endpoint of γ . Observe that every curve in \mathcal{D} that surrounds x must surround y as well, i.e. $\mathcal{D}_x \subseteq \mathcal{D}_y$. Moreover, $|\mathcal{D}_x| = |\mathcal{D}_y| - 1$. By applying the induction hypothesis on x we get an arc γ_x , connecting x to a point at infinity, that intersects every curve in \mathcal{D}_x exactly once and avoids any other curve. By adjoining γ to γ_x , we obtain the desired arc connecting y to a point at infinity. ■

Proof of Theorem 7. We will show that \mathcal{F} can not shatter a set of $k+2$ points. Assume to the contrary that \mathcal{F} shatters a set $S = \{v_1, \dots, v_{k+2}\} \subset P$ of $k+2$ points, i.e. for any subset $V \subseteq S$, there exists a curve $C \in \mathcal{C}$ with $P_C \cap S = V$. For every pair $v_i, v_j \in S$, consider the set of curves $\mathcal{C}_{ij} \subseteq \mathcal{C}$ consisting of all the curves in \mathcal{C} that surround both v_i and v_j . Consider also the set R_{ij} of all the points in the plane which are surrounded by every curve in \mathcal{C}_{ij} . Since \mathcal{C} has the connected intersection property, Lemma 8 implies that R_{ij} is a connected region. Upon drawing an edge between v_i and v_j inside the region R_{ij} , we obtain a drawing of K_{k+2} as a topological graph in the plane which we denote by $\tilde{G} = (S, E)$. We shall investigate the special properties of \tilde{G} , which will eventually lead us to a contradiction.

Claim 11. *Let x be a point in the plane that lies in the unbounded region of $\mathbb{R}^2 \setminus \cup_{C \in \mathcal{C}} C$. Then for every vertex $v_i \in S$ one can draw an arc γ_i , connecting v_i and x , that does not intersect any curve $C \in \mathcal{C}$ with $P_C \cap S = S \setminus \{v_i\}$. Moreover, this drawing can be such that no two arcs γ_i and γ_j cross.*

Proof. Let \mathcal{D} be the subset of \mathcal{C} consisting of all the curves $C \in \mathcal{C}$ with $|P_C \cap S| = k+1$. Since $k \geq 2$ it follows that $|\mathcal{D}| \geq 3$ and that any three discs bounded by curves in \mathcal{D} have a non-empty intersection. Furthermore, because $\mathcal{D} \subseteq \mathcal{C}$, any two discs bounded by curves in \mathcal{D} have a connected intersection. By Lemma 8, there exists a common interior point to all curves in \mathcal{D} . By Lemma 9, the union of any set of discs bounded by curves in \mathcal{D} is simply connected. Thus, for every vertex $v_i \in S$ one can apply Lemma 10 and draw an arc γ_i , connecting v_i with x , such that γ_i avoids any curve $C \in \mathcal{D}$ with $P_C \cap S = S \setminus \{v_i\}$ and crosses any other curve in \mathcal{D} exactly once. From all the possible drawings of such arcs, we pick one with minimum number of intersection points among the γ_i 's. We shall prove that this minimum is 0. Assume otherwise, then there exists a pair of arcs γ_i and γ_j that cross at a point q . We denote by $\gamma_{i,q}$ and $\gamma_{j,q}$ the portions of γ_i and γ_j , respectively, which connect q with x . Both $\gamma_{i,q}$ and $\gamma_{j,q}$ avoid the curves in \mathcal{D} which do not surround q and intersect once the curves in \mathcal{D} which surround q . By swapping the portions $\gamma_{i,q}$ with $\gamma_{j,q}$ and by a small modification of the drawing, we can eliminate the crossing point q and obtain a new drawing of arcs that has one less crossing point. This new drawing still satisfies the property that for every i , each γ_t crosses the curves in \mathcal{D} which surround v_t exactly once and avoids all the other curves in \mathcal{D} . This constitutes a contradiction to the minimality of the number of intersection points among the arcs γ_t in the selected drawing. ■

Let us draw an arc γ_i for every $v_i \in S$ according to Claim 11. Pick an arc, say γ_1 , and define a cyclic order on the arcs γ_i , according to the counterclockwise order in which they reach x , starting with γ_1 . Assume without loss of generality that this order is $\{\gamma_1, \dots, \gamma_{k+2}\}$.

Claim 12. *For every four distinct vertices $v_i, v_j, v_l, v_m \in S$ the edges (v_i, v_j) and (v_l, v_m) in \tilde{G} cross an odd number of times if and only if i and j separate l and m in the natural cyclic order of $\{1, \dots, k+2\}$.*

Proof. We denote by Δ_{ij} the closed curve that is composed by the arcs γ_i, γ_j and the edge (v_i, v_j) in \tilde{G} . We define Δ_{lm} similarly. The curves Δ_{ij} and Δ_{lm} meet at x . Observe that any other intersection point between Δ_{ij} and Δ_{lm} must be an intersection point of the edges (v_i, v_j) and (v_l, v_m) . To see this, recall that in our drawing no two of the arcs $\gamma_1, \dots, \gamma_{k+2}$ cross. Moreover, an arc γ_t connecting v_t to x may cross only those edges of \tilde{G} that are incident to v_t . This is because \mathcal{F} shatters S and therefore there exists a curve $C \in \mathcal{C}$ with $P_C \cap S = S \setminus \{v_t\}$. By the construction of γ_t it avoids $\text{disc}(C)$. Since any edge in \tilde{G} , not incident to v_t , is contained in $\text{disc}(C)$, γ_t cannot cross any edge that is not incident to v_t .

We conclude that any intersection point between Δ_{ij} and Δ_{lm} , other than x , must be an intersection point of the edges (v_i, v_j) and (v_l, v_m) .

If i and j separate l and m in the natural cyclic order $\{1, \dots, k+2\}$, then the curves Δ_{ij} and Δ_{lm} properly *cross* at x . The number of intersection points between two closed curves is even and therefore the edges (v_i, v_j) and (v_l, v_m) must cross an odd number of times.

If i and j do not separate l and m in the natural cyclic order, then Δ_{ij} and Δ_{lm} *touch* at x . As all other intersection points between Δ_{ij} and Δ_{lm} are intersection points of (v_i, v_j) and (v_l, v_m) , it follows that (v_i, v_j) and (v_l, v_m) cross an even number of times. ■

We consider the following two subsets S_1 and S_2 of S :

$$S_1 = \{v_i \in S \mid i \text{ is odd}\} \quad S_2 = \{x_i \in S \mid i \text{ is even}\}.$$

Since \mathcal{F} shatters S , there exist curves $C_1, C_2 \in \mathcal{C}$ such that $P_{C_1} \cap S = S_1$ and $P_{C_2} \cap S = S_2$. We will show that the curves C_1 and C_2 intersect in at least $k+2$ points and obtain a contradiction to the assumption that \mathcal{C} has the k -intersection property.

We call each connected component of $\text{disc}(C_1) \setminus \text{disc}(C_2)$ *an ear*. Similarly, each connected component of $\text{disc}(C_2) \setminus \text{disc}(C_1)$ is called *an ear*. We say that C_1 *enters* C_2 at a crossing point u of C_1 and C_2 if a small enough portion of C_1 that starts at u and continues in the counterclockwise orientation along the curve C_1 is contained in $\text{disc}(C_2)$. Otherwise we say that C_1 *leaves* C_2 at u . We use a similar terminology with respect to C_1 .

Claim 13. *If C_1 and C_2 properly cross in exactly m points, then they create precisely m ears. Moreover, any point that is surrounded by just one curve of C_1 and C_2 must be contained in one of these ears.*

Proof. Let u_1, u_2, \dots, u_m be the set of intersection points of C_1 and C_2 arranged in a counterclockwise order along C_1 . Let w_1, w_2, \dots, w_m be the same set of the intersection points of C_1 and C_2 arranged in a counterclockwise order along C_2 , and assume without loss of generality that $u_1 = w_1$. We first show that $u_i = w_i$ for every $i = 1, \dots, m$. Assume not, then without loss of generality we can assume that $u_2 \neq w_2$ (otherwise, let i be the maximum index such that $u_i = w_i$ and replace u_1 with u_i). Without loss of generality assume that C_2 enters C_1 at u_1 . Then C_1 leaves C_2 at u_1 . We will get a contradiction by showing that $w_2 = u_2$. Assume to the contrary that $w_2 = u_j$ for some $2 < j \leq m$. Then $u_2 = w_l$ for some $2 < l \leq m$. The curve C_1 must enter C_2 at the point $u_2 = w_l$ because it leaves C_2 at u_1 . Therefore, C_2 leaves C_1 at w_l and consequently must enter C_1 at the point w_{l-1} . It follows that the portion δ of C_2 between w_1 and w_2 in the counterclockwise direction along C_2 is contained in $\text{disc}(C_1)$. Similarly, the portion δ' of C_2 between w_{l-1} and w_l in the counterclockwise direction along C_2 is contained in $\text{disc}(C_1)$. δ and δ' split $\text{disc}(C_1)$ into three regions A_1, A_2 , and A_3 , where A_1 is the region bounded by δ and a portion of C_1 , A_2 is the region bounded by both δ and δ' and two portions of C_1 , and A_3 is the region bounded by δ' and a portion of C_1 .

The portion γ of C_1 between $u_1 = w_1$ and $u_2 = w_l$ in the counterclockwise direction along C_1 is connecting a point on δ , namely, w_1 , with a point on δ' , namely, w_l . Since u_1 and u_2 are the only intersection points of C_1 and C_2 on γ , it follows that γ is contained in the boundary of A_2 .

Because C_1 leaves C_2 at u_1 and enters C_2 at u_2 , it must be that γ lies entirely outside of $\text{disc}(C_2)$. It follows that the interior of A_1 must contain points of $\text{disc}(C_1) \cap \text{disc}(C_2)$, and similarly, the interior of A_3 must contain points of $\text{disc}(C_1) \cap \text{disc}(C_2)$. This is a contradiction to the assumption that the interior of $\text{disc}(C_1) \cap \text{disc}(C_2)$ is a connected set. We conclude that $u_i = w_i$ for every $i = 1, \dots, m$.

For every $1 \leq i \leq m$ the portion of C_1 and C_2 between u_i and u_{i+1} forms an ear. Hence, there are at least m ears. We consider $C_1 \cup C_2$ as a planar graph with m vertices and $2m$ edges. By Euler's formula we have $m - 2m + F = 2$, where F is the number of faces created by C_1 and C_2 . Hence, $F = m + 2$. This number includes the unbounded face, namely $\mathbb{R}^2 \setminus (\text{disc}(C_1) \cup \text{disc}(C_2))$, as well as the intersection $\text{disc}(C_1) \cap \text{disc}(C_2)$. We deduce that there are exactly m ears and that every point that is surrounded by one curve but not by the other must lie in one of these ears. ■

We now show that the curves C_1 and C_2 cross in at least $k + 2$ points and thus obtain a contradiction to our assumption that \mathcal{C} has the k -intersection property.




Figure 1: The curves C_1 and C_2

Note that each vertex in S_1 is surrounded by C_1 but not by C_2 . Therefore, by Claim 13, each vertex in S_1 belongs to an ear. Similarly, every vertex in S_2 belongs to an ear. Obviously, a vertex in S_1 and a vertex in S_2 cannot belong to the same ear. More generally, we claim that no two vertices of S may belong to the same ear. Assume to the contrary that $v_i, v_j \in S_1$ belong to the same ear R (we argue similarly, if two vertices in S_2 belong to the same ear). R is contained in $\text{disc}(C_1)$. Draw an arc γ inside R connecting v_i to v_j (see Figure 1). The edge of \tilde{G} connecting v_i and v_j together with γ form a closed curve \tilde{C} that lies inside $\text{disc}(C_1)$. The vertices $v_{i+1}, v_{j+1} \in S_2$ are surrounded by C_2 but not by C_1 and therefore, any arc connecting v_{i+1} and v_{j+1} must cross \tilde{C} an even number of times. The edge of \tilde{G} between v_{i+1} and v_{j+1} crosses the edge of \tilde{G} between v_i and v_j an odd number of times but does not cross γ , as γ lies entirely outside $\text{disc}(C_2)$. Hence, the edge of \tilde{G} connecting v_{i+1} and v_{j+1} crosses \tilde{C} an odd number of times, a contradiction. We conclude that each vertex in S belong to a unique ear. This implies that there are at least $k + 2$ ears. It follows from Claim 13 that C_1 and C_2 intersects in at least $k + 2$ points, which is the desired contradiction.

This also concludes the proof of Theorem 7, as we have shown that \mathcal{F} does not shatter any set of $k + 2$ points. ■

It is an immediate corollary of Theorem 7 and the Perles-Sauer-Shelah theorem that if P

is a set of n points in the plane and \mathcal{C} is a family of simple closed curves with the k -intersection property and the connected intersection property, then $\mathcal{F} = \{P_C \mid C \in \mathcal{C}\}$ consists of $O(n^{k+1})$ members.

We will show, by a construction, that this bound can indeed be attained. For every fixed even number $k \geq 0$, we will construct a set of n points and a family \mathcal{C} of bi-infinite x -monotone curves with the k -intersection property such that the family of all ℓ -sets (for all values of ℓ between 1 and n) of P with respect to \mathcal{C} consists of $\Omega(n^{k+1})$ members. It is then an easy excersize to modify \mathcal{C} to be a family of simple closed curves with the connected intersection property and the k -intersection property, closing each curve at infinity.

Take P to be the integer lattice points $P = \{(a, b) \mid 1 \leq a \leq k+1 \text{ and } 1 \leq b \leq \frac{n}{k+1}\}$. Then for every $(k+1)$ -tuple $(b_1, \dots, b_{k+1}) \in \{1, \dots, \frac{n}{k+1}\}^{k+1}$, let $C_{b_1, \dots, b_{k+1}}$ be the graph of the polynomial of degree at most k passing through each of the points $(i, b_i + \frac{1}{2})$ for $i = 1, \dots, k+1$. Let \mathcal{C} be the collection of all these curves. Because each of the curves in \mathcal{C} is a graph of a polynomial of degree at most k , it follows immediately that \mathcal{C} has the k -intersection property. Observe that each curve in \mathcal{C} defines a unique ℓ -set (for some ℓ). Finally, note that the number of curves in \mathcal{C} is $(\frac{n}{k+1})^{k+1} = \Omega(n^{k+1})$, as required.

References

- [1] P. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir, and S. Smorodinsky, Lenses in Arrangements of Pseudocircles and their Applications, *J. ACM*, **51**, (2004), 139–186.
- [2] N. Alon, G. Kalai, J. Matoušek, and R. Meshulam, Transversal numbers for hypergraphs arising in geometry, *Adv. Appl. Math.*, to appear.
- [3] N. Alon and D. J. Kleitman, Piercing convex sets and the Hadwiger-Debrunner (p, q) -problem, *Adv. Math.* **96** (1992), 103–112.
- [4] T. K. Dey, Improved bounds for planar k -sets and related problems. *Discrete Comput. Geom.* **19** (1998), no. 3, 373–382.
- [5] E. Helly, Über Systeme abgeschlossener Mengen mit gemeinschaftlichen Punkten, *Monatshefte d. Mathematik* **37** (1930), 281–302.
- [6] J. Molnár, Über eine Verallgemeinerung auf die Kugelfläche eines topologischen Satzes von Helly, *Acta Math. Acad. Sci.* **7** (1956), 107–108.
- [7] N. Sauer, On the density of families of sets, *Journal of Combinatorial Theory, Series A* **25** (1972), 80–83.
- [8] S. Shelah, A combinatorial problem, stability and order for models and theories in infinite languages, *Pacific J. Math.* **41** (1972), 247–261.
- [9] J. Snoeyink and J. Hershberger, Sweeping arrangements of curves, *DIMACS Series in Discrete Mathematics, Discrete and Computational Geometry, the DIMACS Special Year 6* (1991), 309–349.
- [10] G. Tóth, Point sets with many k -sets, *Discrete Comput. Geom.* **26** (2001) no. 2, 187–194.
- [11] V.N. Vapnik and A. Ya. Chervonenkis, On the uniform convergence of relative frequencies of events to their probabilities, *Theory Probab. Appl.* **16** (1971), 264–280.