

On the Zarankiewicz Problem for Intersection Hypergraphs

Nabil H. Mustafa* and János Pach**

Abstract. Let d and t be fixed positive integers, and let $K_{t,\dots,t}^d$ denote the complete d -partite hypergraph with t vertices in each of its parts, whose hyperedges are the d -tuples of the vertex set with precisely one element from each part. According to a fundamental theorem of extremal hypergraph theory, due to Erdős [7], the number of hyperedges of a d -uniform hypergraph on n vertices that does not contain $K_{t,\dots,t}^d$ as a subhypergraph, is $n^{d-\frac{1}{t^{d-1}}}$. This bound is not far from being optimal.

We address the same problem restricted to *intersection hypergraphs* of $(d-1)$ -dimensional *simplices* in \mathbb{R}^d . Given an n -element set \mathcal{S} of such simplices, let $\mathcal{H}^d(\mathcal{S})$ denote the d -uniform hypergraph whose vertices are the elements of \mathcal{S} , and a d -tuple is a hyperedge if and only if the corresponding simplices have a point in common. We prove that if $\mathcal{H}^d(\mathcal{S})$ does not contain $K_{t,\dots,t}^d$ as a subhypergraph, then its number of edges is $O(n)$ if $d = 2$, and $O(n^{d-1+\epsilon})$ for any $\epsilon > 0$ if $d \geq 3$. This is almost a factor of n better than Erdős's above bound. Our result is tight, apart from the error term ϵ in the exponent.

In particular, for $d = 2$, we obtain a theorem of Fox and Pach [8], which states that every $K_{t,t}$ -free intersection graph of n *segments* in the plane has $O(n)$ edges. The original proof was based on a separator theorem that does not generalize to higher dimensions. The new proof works in any dimension and is simpler: it uses *size-sensitive cuttings*, a variant of random sampling. We demonstrate the flexibility of this technique by extending the proof of the planar version of the theorem to intersection graphs of x -monotone curves.

1 Introduction

Let \mathcal{H} be a d -uniform hypergraph on n vertices. One of the fundamental questions of extremal graph and hypergraph theory goes back to Turán and Zarankiewicz: What is the largest number $\text{ex}^d(n, K)$ of hyperedges (or, in short, edges) that \mathcal{H} can have if it contains no subhypergraph isomorphic to fixed d -uniform hypergraph K . (See Bollobás [2]). In the most applicable special case, $K = K^d(t, \dots, t)$ is the complete d -partite hypergraph on the vertex set V_1, \dots, V_d with $|V_1| = \dots = |V_d| = t$, consisting of all d -tuples that contain one point from each V_i . For graphs ($d = 2$), it was proved

* Université Paris-Est, Laboratoire d'Informatique Gaspard-Monge, Equipe A3SI, ESIEE Paris. Supported by the grant ANR SAGA (JCJC-14-CE25-0016-01). E-mail: mustafan@esiee.fr.

** EPFL, Lausanne and Rényi Institute, Budapest. Supported by Swiss National Science Foundation Grants 200020-144531 and 200021-137574. E-mail: pach@cims.nyu.edu.

by Erdős (1938) and Kővári-Sós-Turán [14] that

$$\text{ex}^2(n, K_{t,t}^2) \leq n^{2-1/t}.$$

The order of magnitude of this estimate is known to be best possible only for $t = 2$ and 3 (Reiman [18]; Brown [3]). The constructions for which equality is attained are algebraic.

Erdős (1964) generalized the above statement to d -uniform hypergraphs for all $d \geq 2$. A hypergraph \mathcal{H} is said to be K -free if it contains no copy of K as a (not necessarily induced) subhypergraph.

Theorem A ([7]). *The maximum number of hyperedges that a d -uniform, $K_{t,\dots,t}^d$ -free hypergraph of n vertices can have satisfies*

$$\text{ex}^d(n, K_{t,\dots,t}^d) \leq n^{d-\frac{1}{t^{d-1}}}.$$

It was further shown in [7] that this bound cannot be substantially improved. There exists an absolute constant $C > 0$ (independent of n, t, d) such that

$$\text{ex}^d(n, K_{t,\dots,t}^d) \geq n^{d-\frac{C}{t^{d-1}}}.$$

In particular, for every $\epsilon > 0$, there exist $K_{t,\dots,t}^d$ -free d -uniform hypergraphs with $t \approx (C/\epsilon)^{1/(d-1)}$, having at least $n^{d-\epsilon}$ edges. The construction uses the probabilistic method.

In certain geometric scenarios, better bounds are known. For instance, consider a bipartite graph with $2n$ vertices that correspond to n distinct points and n distinct lines in the plane, and a vertex representing a point p is connected to a vertex representing a line ℓ if and only if p is incident to ℓ (i.e., $p \in \ell$). Obviously, this graph is $K_{2,2}^2$ -free (in short, $K_{2,2}$ -free). Therefore, by the above result, it has at most $O(n^{3/2})$ edges. On the other hand, according to a celebrated theorem of Szemerédi and Trotter (1983), the number of edges is at most $O(n^{4/3})$, and the order of magnitude of this bound cannot be improved. In [15], this result was generalized to incidence graphs between points and more complicated curves in the plane.

Given a set \mathcal{S} of geometric objects, their *intersection graph* $\mathcal{H}(\mathcal{S})$ is defined as a graph on the vertex set \mathcal{S} , in which two vertices are joined by an edge if and only if the corresponding elements of \mathcal{S} have a point in common. To better understand the possible intersection patterns of edges of a *geometric graph*, that is, of a graph drawn in the plane with possibly crossing straight-line edges, Pach and Sharir [15] initiated the investigation of the following problem. What is the maximum number of edges in a $K_{t,t}$ -free intersection graph of n segments in the plane? The Kővári-Sós-Turán theorem (Theorem A for $d = 2$) immediately implies the upper bound $n^{2-1/t}$. Pach and Sharir managed to improve this bound to $O(n)$ for $t = 2$ and to $O(n \log n)$ for any larger (but fixed) value of t . They conjectured that here $O(n \log n)$ can be replaced by $O(n)$ for

every t , which was proved by Fox and Pach [8]. They later extended their proof to *string graphs*, that is, to intersection graphs of arbitrary continuous arcs in the plane [9, 10]. Some weaker results were established by Radoičić and Tóth [17] by the “discharging method”.

The aim of the present note is to generalize the above results to *d-uniform intersection hypergraphs of (d − 1)-dimensional simplices* in \mathbb{R}^d , for any $d \geq 2$. The arguments used in the above papers are based on planar separator theorems that do not seem to allow higher dimensional extensions applicable to our problem.

Given an n -element set \mathcal{S} of $(d - 1)$ -dimensional simplices in general position in \mathbb{R}^d , let $\mathcal{H}^d(\mathcal{S})$ denote the d -uniform hypergraph on the vertex set \mathcal{S} , consisting of all unordered d -tuples of elements $\{S_1, \dots, S_d\} \subset \mathcal{S}$ with $S_1 \cap \dots \cap S_d \neq \emptyset$. We prove the following theorem, providing an upper bound on the number of hyperedges of a $K_{t, \dots, t}^d$ -free intersection hypergraph $\mathcal{H}^d(\mathcal{S})$ of $(d - 1)$ -dimensional simplices. This bound is almost a factor of n better than what we obtain using the abstract combinatorial bound of Erdős (Theorem A), and it does not depend strongly on t .

Theorem 1. *Let $d, t \geq 2$ be integers, let \mathcal{S} be an n -element set of $(d - 1)$ -dimensional simplices in \mathbb{R}^d , and let $\mathcal{H}^d(\mathcal{S})$ denote its d -uniform intersection hypergraph.*

If $\mathcal{H}^d(\mathcal{S})$ is $K_{t, \dots, t}^d$ -free, then its number of edges is $O(n^{d-1+\epsilon})$ for any $\epsilon > 0$. For $d = 2$, the number of edges is at most $O(n)$.

To see that this bound is nearly optimal for every d , fix a hyperplane h in \mathbb{R}^d with normal vector v , and pick $d - 1$ sets P_1, \dots, P_{d-1} , each consisting of $\frac{n-1}{d-1}$ parallel $(d - 2)$ -dimensional planes in h , with the property that any $d - 1$ members of $P_1 \cup \dots \cup P_{d-1}$ that belong to different P_i s have a point in common. For each i ($1 \leq i \leq d - 1$), replace every plane $p_i \in P_i$ by a hyperplane h_i parallel to u such that $h_i \cap h = p_i$. Clearly, the d -uniform intersection hypergraph of these n hyperplanes, including h , is $K_{2, \dots, 2}^d$ -free, and its number of edges is $(\frac{n-1}{d-1})^{d-1} = \Omega(n^{d-1})$. In each hyperplane, we can take a large $(d - 1)$ -dimensional simplex so that the intersection pattern of these simplices is precisely the same as the intersection pattern (i.e., the d -uniform intersection hypergraph) of the underlying hyperplanes.

The proof of Theorem 1 is based on a partitioning scheme, which was first formulated by Pellegrini [16]. Given an n -element set \mathcal{S} of $(d - 1)$ -dimensional simplices (or other geometric objects) in \mathbb{R}^d , let m denote the number of hyperedges in their d -uniform intersection hypergraph, that is, the number of d -tuples of elements of \mathcal{S} having a point in common. For a parameter $r \leq n$, a $\frac{1}{r}$ -cutting with respect to \mathcal{S} is a partition of \mathbb{R}^d into simplices such that the interior of every simplex intersects at most n/r elements of \mathcal{S} . The *size* of a $\frac{1}{r}$ -cutting is the number of simplices it consists of. (See Matoušek [13].)

Theorem B ([6, 16]). *Let $d \geq 2$ be an integer, let \mathcal{S} be an n -element set of $(d - 1)$ -dimensional simplices in \mathbb{R}^d , and let m denote the number of d -tuples of simplices in \mathcal{S} having a point in common.*

Then, for any $\epsilon > 0$ and any $r \leq n$, there is a $\frac{1}{r}$ -cutting with respect to \mathcal{S} of size at most

$$C_2 \cdot \left(r + \frac{mr^2}{n^2} \right) \quad \text{if } d = 2, \text{ and}$$

$$C_{d,\epsilon} \cdot \left(r^{d-1+\epsilon} + \frac{mr^d}{n^d} \right) \quad \text{if } d \geq 3.$$

Here C_2 is an absolute constant and $C_{d,\epsilon}$ depends only on d and ϵ .

To construct a $\frac{1}{r}$ -cutting, we have to take a *random sample* $\mathcal{R} \subseteq \mathcal{S}$, where each element of \mathcal{S} is selected with probability r/n . It can be shown that, for every $k > 0$, the expected value of the total number of k -dimensional faces of all cells of the cell decomposition induced by the elements of \mathcal{R} is $O(r^{d-1})$, while the expected number of vertices (0-dimensional faces) is $m(r/n)^d$. This cell decomposition can be further subdivided to obtain a partition of \mathbb{R}^d into simplices that meet the requirements. The expected number of elements of \mathcal{S} that intersect a given cell is at most n/r .

Cuttings have been successfully used before, e.g., for an alternative proof of the Szemerédi-Trotter theorem [5]. In our case, the use of this technique is somewhat un-intuitive, as the size of the cuttings we construct depends on the number of intersecting d -tuples, that is, on the parameter we want to bound. This sets up an unusual recurrence relation, where the required parameter appears on both sides, but nonetheless, whose solution implies Theorem 1.

The use of cuttings is versatile. In particular, we show how this technique can be applied to establish the generalization of Theorem 1 to x -monotone curves with a bounded number of pairwise intersections.

Theorem 2. *Let $\mathcal{S} = \{S_1, \dots, S_n\}$ be a set of n x -monotone curves in \mathbb{R}^2 , and where every pair of curves intersect at most a constant number of times. Let $\mathcal{H}(\mathcal{S})$ denote its intersection graph, and $t \geq 2$ be an integer. If $\mathcal{H}(\mathcal{S})$ is $K_{t,t}^2$ -free, then its number of edges is $O(n)$.*

Fox and Pach [9, 10] managed to prove the same result for all continuous curves in the plane, using separator theorems. The weakness of their method is that it is inherently planar. As we will see, the weakness of the cutting technique is that the cell decomposition defined by the randomly selected objects needs to be further refined. To make sure that the number of cells remains under control in this step, we have to bound the number of d -wise intersection points between the objects.

2 Proof of Theorem 1

For any element $S_k \in \mathcal{S}$, let $\text{supp}S_k$ denote the supporting hyperplane of S_k . By slightly perturbing the arrangements, if necessary, we can assume without loss of gen-

erality that the supporting hyperplanes of the elements of \mathcal{S} are in *general position*, that is,

- (a) no $d - j + 1$ of them have a j -dimensional intersection ($0 \leq j \leq d - 1$), and
- (b) the intersection of any d of them is empty or a point that lies in the relative interior of these d elements.

Theorem 1 is an immediate corollary of the following lemma.

Lemma 1. *Let $d, t \geq 2$ be fixed integers. Let \mathcal{S} be an n -element set of $(d - 1)$ -dimensional simplices in general position in \mathbb{R}^d . Assume that their d -uniform intersection hypergraph $\mathcal{H}^d(\mathcal{S})$ has m edges and is $K_{t, \dots, t}^d$ -free.*

If, for suitable constants $C \geq 1$ and u , there exists a $\frac{1}{r}$ -cutting of size at most $C(r^u + \frac{mr^d}{n^d})$ with respect to \mathcal{S} , consisting of full-dimensional simplices, then

$$m \leq C' \cdot n^u,$$

where C' is another constant (depending on d, t, C , and u).

Proof. For some value of the parameter r to be specified later, construct a $\frac{1}{r}$ -cutting $\{\Delta_1, \dots, \Delta_k\}$ with respect to \mathcal{S} , where $k \leq C(r^u + \frac{mr^d}{n^d})$. Using our assumption that the elements of \mathcal{S} are in general position, we can suppose that all cells Δ_i are full-dimensional.

For every i ($1 \leq i \leq k$), let $\mathcal{S}_i^{\text{int}} \subseteq \mathcal{S}$ denote the set of all elements in \mathcal{S} that intersect the interior of Δ_i . As $\{\Delta_1, \dots, \Delta_k\}$ is a cutting with respect to \mathcal{S} , we have $|\mathcal{S}_i^{\text{int}}| \leq n/r$. Let $\mathcal{S}_i^{\text{bd}} \subseteq \mathcal{S}$ be the set of all elements $S_k \in \mathcal{S}$ such that the supporting hyperplane $\text{supp}S_k$ of S_k contains a j -dimensional face of Δ_i for some j ($0 \leq j \leq d - 1$). Set $\mathcal{S}_i = \mathcal{S}_i^{\text{int}} \cup \mathcal{S}_i^{\text{bd}}$.

Using the general position assumption, we obtain that every j -dimensional face F of Δ_i is contained in the supporting hyperplanes of at most $d - j \leq d$ elements of \mathcal{S} ($0 \leq j \leq d - 1$). The total number of proper faces of Δ_i of all dimensions is smaller than 2^{d+1} , and each is contained in at most d elements of \mathcal{S} . Therefore, we have

$$|\mathcal{S}_i| < |\mathcal{S}_i^{\text{int}}| + |\mathcal{S}_i^{\text{bd}}| \leq n/r + d2^{d+1}.$$

Fix an intersection point $q = S_1 \cap \dots \cap S_d$, and let $\mathcal{S}(q) = \{S_1, \dots, S_d\} \subseteq \mathcal{S}$. Then either

1. q lies in the interior of some Δ_i , in which case $\mathcal{S}(q) \subseteq \mathcal{S}_i^{\text{int}} \subseteq \mathcal{S}_i$, or
2. q lies at a vertex (0-dimensional face) F or in the interior of a j -dimensional face F of some Δ_i , where $1 \leq j \leq d - 1$. Take any $S_k \in \mathcal{S}(q)$. If $F \subset \text{supp}S_k$, then $S_k \in \mathcal{S}_i^{\text{bd}} \subseteq \mathcal{S}_i$. If $F \not\subset \text{supp}S_k$, then S_k intersects the interior of Δ_i , and since q lies in the relative interior of S_k , we have that $S_k \in \mathcal{S}_i^{\text{int}} \subseteq \mathcal{S}_i$.

In both cases, $\mathcal{S}(q) \subseteq \mathcal{S}_i$.

This means that if for each i we bound the number of d -wise intersection points between the elements of \mathcal{S}_i , and we add up these numbers, we obtain an upper bound on $\mathcal{H}^d(\mathcal{S})$, the number of d -wise intersection points between the elements of \mathcal{S} . Within each \mathcal{S}_i , we apply the abstract hypergraph-theoretic bound of Erdős (Theorem A) to conclude that $\mathcal{H}^d(\mathcal{S}_i)$ has at most $|\mathcal{S}_i|^{d-1/t^{d-1}}$ edges. Hence,

$$m \leq \sum_{i=1}^k |\mathcal{S}_i|^{d-1/t^{d-1}}.$$

As $|\mathcal{S}_i| \leq n/r + d2^{d+1}$, substituting the bound on k , we get

$$\begin{aligned} m &\leq C \cdot \left(r^u + \frac{mr^d}{n^d}\right) \cdot \left(d2^{d+1} + \frac{n}{r}\right)^{d-1/t^{d-1}} \\ &\leq 2^d C \cdot \left(r^u + \frac{mr^d}{n^d}\right) \cdot \left(\frac{n}{r}\right)^{d-1/t^{d-1}}, \end{aligned}$$

provided that $\frac{n}{r} \geq d2^{d+1}$. Setting $r = \frac{n}{C_0}$, where $C_0 = (2^{d+1}C)^{t^{d-1}}$, we obtain

$$\begin{aligned} m &\leq 2^d C \left(\frac{n^u}{C_0^u} + \frac{m}{C_0^d}\right) C_0^{d-1/t^{d-1}} \\ m &\leq \frac{C_0^{d-u}}{2} n^u + \frac{m}{2}, \end{aligned}$$

which implies that $m \leq C_0^{d-u} n^u$, as required. \square

Now Theorem 1 follows from Theorem B, as one can choose $u = 1$ if $d = 2$ and $u = d - 1 + \epsilon$ if $d \geq 3$. This completes the proof.

Remark 1. Every d -uniform hypergraph \mathcal{H} has a d -partite subhypergraph \mathcal{H}' that has at least $\frac{d!}{d^d}$ times as many hyperedges as \mathcal{H} . Therefore, if K is d -partite, the maximum number of hyperedges that a d -partite K -free hypergraph on n vertices can have is within a factor of $\frac{d!}{d^d}$ from the same quantity over all K -free hypergraphs on n vertices. If instead of abstract hypergraphs, we restrict our attention to intersection graphs or hypergraphs of geometric objects, the order of magnitudes of the two functions may substantially differ. Given two sets of segments \mathcal{S} and \mathcal{T} in the plane with $|\mathcal{S}| = |\mathcal{T}| = n$, let $\mathcal{B}(\mathcal{S}, \mathcal{T})$ denote their *bipartite* intersection graph, in which the vertices representing \mathcal{S} and \mathcal{T} form two independent sets, and a vertex representing a segment in \mathcal{S} is joined to a vertex representing a segment in \mathcal{T} are joined by an edge if and only if they intersect. It was shown in [11] that any $K_{2,2}$ -free bipartite intersection graph of n vertices has $O(n^{4/3})$ edges and that this bound is tight. In fact, this result generalizes the Szemerédi-Trotter theorem mentioned above. On the other hand, if we assume that the (non-bipartite) intersection graph associated with the set $\mathcal{S} \cup \mathcal{T}$, which contains the bipartite graph $\mathcal{B}(\mathcal{S}, \mathcal{T})$, is also $K_{2,2}$ -free, then Theorem 1 implies that the number of

edges drops to linear in n . In the examples where $\mathcal{B}(\mathcal{S}, \mathcal{T})$ has a superlinear number of edges, there must be many intersecting pairs of segments in \mathcal{S} or in \mathcal{T} .

Remark 2. The key assumption in Lemma 1 is that there exists a $\frac{1}{r}$ -cutting, whose size is sensitive to the number of intersecting d -tuples of objects. Under these circumstances, in terms of the smallest size of a $\frac{1}{r}$ -cutting, one can give an upper bound on the number of edges of $K_{t,\dots,t}^d$ -free intersection hypergraphs with n vertices. For $d = 3$, we know some stronger bounds on the size of vertical decompositions of space induced by a set of triangles [6, 19], which imply the existence of $\frac{1}{r}$ -cuttings of size $O(r^2\alpha(r) + \frac{mr^3}{n^3})$. Thus, in this case, we can deduce from Lemma 1 that every 3-uniform $K_{t,t,t}^3$ -free intersection hypergraph of n triangles has $O(n^2\alpha(n))$ edges. It is an interesting open problem to establish nearly tight bounds on the maximum number of edges that a d -uniform $K_{t,\dots,t}^d$ -free intersection hypergraph induced by n semialgebraic sets in \mathbb{R}^d can have. Lemma 1 does not apply in this case, because in the best currently known constructions of cuttings for semialgebraic sets, the exponent u is larger than d [1].

3 Proof of Theorem 2

Given a set \mathcal{S} of curves in \mathbb{R}^2 , we will assume that no three curves pass through a common point, and that no two intersection points have the same y -coordinate.

The proof of Theorem 2 follows from an appropriate modification of cuttings for x -monotone curves. Define a *cell* to be a closed set in \mathbb{R}^2 that is homeomorphic to a disk. A cell c is *induced by* \mathcal{S} if its boundary is composed of subcurves of elements of \mathcal{S} , and (possibly) line segments. The size of a cell c is the number of its boundary curves and segments. For a cell c and a set $\mathcal{R} \subseteq \mathcal{S}$, define $\mathcal{R}_c^{\text{int}}$ to be the curves in \mathcal{R} intersecting the interior of c . A *decomposition* \mathcal{T} of a set $\mathcal{R} \subseteq \mathcal{S}$ is a set of interior-disjoint cells induced by \mathcal{R} , and covering \mathbb{R}^2 such that $\mathcal{R}_c^{\text{int}} = \emptyset$ for each cell $c \in \mathcal{T}$. The size of a decomposition is its number of cells. Each cell c will be associated with a unique subset $\mathcal{S}_c^{\text{bd}} \subseteq \mathcal{S}$. Given \mathcal{S} , the set of decompositions of every $\mathcal{R} \subseteq \mathcal{S}$ is called a *canonical decomposition scheme* if it satisfies the following two properties (see [13, Section 6.5] for details):

1. for every cell c in all decompositions, the size of $\mathcal{S}_c^{\text{bd}}$ can be bounded from above by a constant, and
2. a cell c belongs to the decomposition of \mathcal{R} if and only if $\mathcal{S}_c^{\text{bd}} \subseteq \mathcal{R}$ and $\mathcal{R} \cap \mathcal{S}_c^{\text{int}} = \emptyset$.

A $(1/r)$ -cutting \mathcal{H} for \mathcal{S} is a partition of \mathbb{R}^2 into interior-disjoint cells of bounded size such that the interior of each cell $c \in \mathcal{H}$ is intersected by at most n/r curves of \mathcal{S} . Given a canonical decomposition scheme for a set of objects, the existence of small-sized cuttings follows from

Theorem C ([6]). *Let \mathcal{S} be a set of n objects in \mathbb{R}^2 , and let $r \leq n$ be a parameter. Assume that there exists a $(1/a)$ -cutting for any $\mathcal{S}' \subseteq \mathcal{S}$ and any $a > 0$, of size $O(a^C)$ where C is some constant. Then there exists a $(1/r)$ -cutting for \mathcal{S} of size $O(\tau(r))$,*

where $\tau(r)$ is the expected number of cells in the canonical decomposition of a random subset of \mathcal{S} where each element of \mathcal{S} is picked with probability r/n .

From here, using standard methods, one can deduce the following theorem, whose proof is sketched here, for completeness.

Theorem 3 (Cuttings for x -monotone curves). *Let \mathcal{S} be an n -element set of x -monotone curves in \mathbb{R}^2 , such that every pair of curves intersect at most a constant number of times. Let m be the number of pairs of intersecting curves in \mathcal{S} . Then there exists a $(1/r)$ -cutting for \mathcal{S} of size $O(r + mr^2/n^2)$.*

Proof. The canonical decomposition of any $\mathcal{S}' \subseteq \mathcal{S}$ will be the *vertical decomposition* of \mathcal{S}' . Recall that the vertical decomposition of \mathcal{S}' is constructed by extending, from each endpoint of a curve of \mathcal{S}' as well as from each intersection point, a vertical segment above and below until it hits another curve of \mathcal{S}' . It was verified in [4] that this decomposition is a canonical decomposition scheme.

The vertical decomposition of any $\mathcal{S}' \subseteq \mathcal{S}$ has size $O(|\mathcal{S}'| + I_{\mathcal{S}'})$, where $I_{\mathcal{S}'}$ is the number of intersection points of the curves in \mathcal{S}' . Let $m_{\mathcal{S}'}$ be the number of edges in the intersection graph of \mathcal{S}' . As each pair of curves can intersect at most a constant number of times, we have $I_{\mathcal{S}'} = O(m_{\mathcal{S}'})$. Thus, the vertical decomposition of \mathcal{S}' has size $O(|\mathcal{S}'| + m_{\mathcal{S}'})$.

Let $\mathcal{R} \subseteq \mathcal{S}$ be a set formed by picking each curve of \mathcal{S} with probability $p = r/n$. The expected number of edges in the intersection graph of \mathcal{R} is $\mathbf{E}[m_{\mathcal{R}}] = mp^2$. Therefore, the expected size of the canonical decomposition of \mathcal{R} is $O(|\mathcal{S}|p + mp^2) = O(r + mr^2/n^2)$.

It remains to show that for any $a > 0$ and for any $\mathcal{S}' \subseteq \mathcal{S}$, there exists a $(1/a)$ -cutting for \mathcal{S}' of size $O(a^C)$. Fix a set \mathcal{S}' , and let \mathcal{U} be the set of all cells present in the vertical decomposition of any subset \mathcal{S}'' of \mathcal{S}' . Construct the following set-system on \mathcal{S}' :

$$\Phi(\mathcal{S}') = \{\mathcal{S}' \cap \text{interior}(U) \mid U \in \mathcal{U}\}$$

Each boundary vertex of a cell in \mathcal{U} is an (i) intersection point of two curves in \mathcal{S}' , or (ii) the endpoint of a curve in \mathcal{S}' , or (iii) the intersection point of a curve in \mathcal{S}' and a vertical line passing through either an endpoint or an intersection point of two curves in \mathcal{S}' . Since every pair of curves intersect at most a constant number of times, the total number of such vertices is $O(|\mathcal{S}'|^3)$. Each cell in the vertical decomposition of any subset of \mathcal{S}' can be uniquely identified by the sequence of its boundary vertices. As each cell has at most 4 boundary vertices, we get that $|\mathcal{U}| = O(|\mathcal{S}'|^{12})$. Therefore, $|\Phi(\mathcal{S}')| = O(|\mathcal{S}'|^{12})$, which implies that the VC-dimension of $\Phi(\mathcal{S}')$ is bounded by a constant. By the ϵ -net theorem of Haussler and Welzl [12], there exists a subset $\mathcal{R}' \subseteq \mathcal{S}'$ of size $O(a \log a)$, such that $\mathcal{S}'' \cap \mathcal{R}' \neq \emptyset$ for any $\mathcal{S}'' \in \Phi(\mathcal{S}')$ of size at least $|\mathcal{S}'|/a$. The required $(1/a)$ -cutting is the vertical decomposition of \mathcal{R}' , as by property 2 of canonical decompositions, for any cell c in the vertical decomposition, we have $|\mathcal{S}_c^{\text{int}}| < |\mathcal{S}'|/a$. Finally, note that the size of the vertical decomposition of \mathcal{R}' is $O(a^2 \log^2 a)$. \square

Now are in a position to complete the proof of Theorem 2. Construct a $(1/r)$ -cutting Π for \mathcal{S} , of size $C_3 \cdot (r + mr^2/n^2)$, where C_3 is a constant. For each cell $c \in \Pi$, let $\mathcal{S}_c = \mathcal{S}_c^{\text{bd}} \cup \mathcal{S}_c^{\text{int}}$. Note that $|\mathcal{S}_c| \leq O(1) + n/r$ for all c . For a point $p \in S_i \cap S_j$, there are three possibilities:

1. p lies in the interior of a cell $c \in \Pi$. Then $S_i, S_j \in \mathcal{S}_c^{\text{int}}$.
2. p lies in the interior of the boundary of a cell c . By the assumption that every pair of curves of \mathcal{S} intersect in at most a constant number of points and no three pass through a common point, $S_i \in \mathcal{S}_c^{\text{bd}}$ and $S_j \in \mathcal{S}_c^{\text{int}}$ (or vice versa).
3. p lies at a vertex of c . The total number of such intersection points is bounded by the number of vertices of Π , $O(r + mr^2/n^2)$.

This implies that

$$m \leq O\left(r + \frac{mr^2}{n^2}\right) + \sum_{c \in \Pi} |\mathcal{S}_c|^{2-1/t}$$

As in proof of Theorem 1, set $r = n/C$ for a sufficiently large constant C to get $m = O(n)$.

Remark. Theorem 2 implies that if the intersection graph of n constant-degree algebraic curves in the plane is $K_{t,t}$ -free, then it has $O(n)$ edges.

References

1. P. K. Agarwal, J. Matousek, and M. Sharir. On range searching with semialgebraic sets. II. *SIAM J. Comput.*, 42(6):2039–2062, 2013.
2. B. Bollobás. *Modern Graph Theory*. Springer-Verlag, 1998.
3. W. G. Brown. On graphs that do not contain a Thomsen graph. *Canad. Math. Bull.*, 9:281285, 1966.
4. B. Chazelle, H. Edelsbrunner, L. J. Guibas, M. Sharir, and J. Snoeyink. Computing a face in an arrangement of line segments and related problems. *SIAM J. Comput.*, 22(6):1286–1302, 1993.
5. K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir, and E. Welzl. Combinatorial complexity bounds for arrangements of curves and spheres. *Discrete & Computational Geometry*, 5:99–160, 1990.
6. M. de Berg and O. Schwarzkopf. Cuttings and applications. *Int. J. Comput. Geometry Appl.*, 5(4):343–355, 1995.
7. P. Erdős. On extremal problems of graphs and generalized graphs. *Israel J. Math.*, 2:183–190, 1964.
8. J. Fox and J. Pach. Separator theorems and Turán-type results for planar intersection graphs. *Advances in Mathematics*, 219(3):1070–1080, 2008.
9. J. Fox and J. Pach. A separator theorem for string graphs and its applications. *Comb. Probab. Comput.*, 19(3):371–390, 2010.
10. J. Fox and J. Pach. Applications of a new separator theorem for string graphs. *Combinatorics, Probability and Computing*, 23:66–74, 2014.
11. J. Fox, J. Pach, A. Sheffer, A. Suk, and J. Zahl. A semi-algebraic version of Zarankiewicz’s problem. *ArXiv e-prints*, 2014.

12. D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.
13. J. Matoušek. *Lectures in Discrete Geometry*. Springer-Verlag, New York, NY, 2002.
14. T. Kővári, V. T. Sós, and P. Turán. On a problem of K. Zarankiewicz. *Colloquium Math.*, 3:50–57, 1954.
15. J. Pach and M. Sharir. On planar intersection graphs with forbidden subgraphs. *J. Graph Theory*, 59(3):205–214, 2008.
16. M. Pellegrini. On counting pairs of intersecting segments and off-line triangle range searching. *Algorithmica*, 17(4):380–398, 1997.
17. R. Radoičić and G. Tóth. The discharging method in combinatorial geometry and the Pach–Sharir conjecture. In J. E. Goodman, J. Pach, and J. Pollack, editors, *Proceedings of the Joint Summer Research Conference on Discrete and Computational Geometry*, volume 453, pages 319–342. Contemporary Mathematics, AMS, 2008.
18. I. Reiman. Über ein problem von K. Zarankiewicz. *Acta Mathematica Academiae Scientiarum Hungarica*, 9:269–273, 1958.
19. B. Tagansky. A new technique for analyzing substructures in arrangements of piecewise linear surfaces. *Discrete & Computational Geometry*, 16(4):455–479, 1996.