Zarankiewicz's Problem via ϵ-t-Nets

Chaya Keller
Department of Computer Science, Ariel University, Israel

Shakhar Smorodinsky
Department of Computer Science, Ben-Gurion University of the NEGEV, Beer Sheva, Israel

Abstract

The classical Zarankiewicz's problem asks for the maximum number of edges in a bipartite graph on n vertices which does not contain the complete bipartite graph $K_{t,t}$. Kővári, Sós and Turán proved an upper bound of $O(n^{2-\frac{1}{t}})$. Fox et al. obtained an improved bound of $O(n^{2-\frac{1}{d}})$ for graphs of VC-dimension d (where $d < t$). Basit, Chernikov, Starchenko, Tao and Tran improved the bound for the case of semilinear graphs. Chan and Har-Peled further improved Basit et al.'s bounds and presented (quasi-)linear upper bounds for several classes of geometrically-defined incidence graphs, including a bound of $O(n \log\log n)$ for the incidence graph of points and pseudo-discs in the plane.

In this paper we present a new approach to Zarankiewicz's problem, via ϵ-t-nets – a recently introduced generalization of the classical notion of ϵ-nets. Using the new approach, we obtain a sharp bound of $O(n)$ for the intersection graph of two families of pseudo-discs, thus both improving and generalizing the result of Chan and Har-Peled from incidence graphs to intersection graphs. We also obtain a short proof of the $O(n^{2-\frac{1}{d}})$ bound of Fox et al., and show improved bounds for several other classes of geometric intersection graphs, including a sharp $O(n\frac{\log n}{\log\log n})$ bound for the intersection graph of two families of axis-parallel rectangles.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Zarankiewicz's Problem, ϵ-t-nets, pseudo-discs, VC-dimension

Digital Object Identifier 10.4230/LIPIcs.SoCG.2024.66

Related Version *Full Version*: `https://arxiv.org/pdf/2311.13662.pdf`

Funding *Chaya Keller*: Research partially supported by Grant 1065/20 from the Israel Science Foundation.
Shakhar Smorodinsky: Research partially supported by Grant 1065/20 from the Israel Science Foundation.

1 Introduction

Zarankiewicz's problem. A central research area in extremal combinatorics is *Turán-type questions*, which ask for the maximum number of edges in a graph on n vertices that does not contain a copy of a fixed graph H. This question was raised in 1941 by Turán, who showed that the maximum number of edges in a K_r-free graph on n vertices is $(1 - \frac{1}{r-1} + o(1))\frac{n^2}{2}$. Soon after, Erdős, Stone and Simonovits solved the problem for all non-bipartite graphs H. They showed that the maximum number is $(1 - \frac{1}{\chi(H)-1} + o(1))\frac{n^2}{2}$, where $\chi(H)$ is the chromatic number of H.

The bipartite case turned out to be significantly harder, and the question is still widely open for most bipartite graphs (see the survey [29]). The case of H being a complete bipartite graph was first studied by Zarankiewicz in 1951:

▶ **Problem 1** (Zarankiewicz's problem). *What is the maximum number of edges in a $K_{t,t}$-free bipartite graph on n vertices?*

40th International Symposium on Computational Geometry (SoCG 2024).
Editors: Wolfgang Mulzer and Jeff M. Phillips; Article No. 66; pp. 66:1–66:15
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In one of the cornerstone results of extremal graph theory, Kővári, Sós and Turán [25] proved an upper bound of $O(n^{2-\frac{1}{t}})$. This bound is sharp for $t = 2, 3$. For $t = 2$, a tightness example is the incidence graph of points and lines in a finite projective plane. Similarly, a tightness example for $t = 3$ is an incidence graph of points and spheres (of a carefully chosen fixed radius) in a three-dimensional finite affine space (see [9]). The question whether the Kővári-Sós-Turán theorem is tight for $t \geq 4$ is one of the central open problems in extremal graph theory.

VC-dimension and Fox et al.'s bound for Zarankiewicz's problem. The VC-dimension of a hypergraph is a measure of its complexity, which plays a central role in statistical learning, computational geometry, and other areas of computer science and combinatorics (see, e.g., [4, 8, 28]). The VC-dimension of a hypergraph $H = (V, \mathcal{E})$ is the largest integer d for which there exists $S \subset V, |S| = d$, such that for every subset $B \subset S$, one can find a hyperedge $e \in \mathcal{E}$ with $e \cap S = B$. Such a set S is said to be *shattered*. The *primal shatter function* of H is $\pi_H(m) = \max_{S \subset V, |S|=m} |\{S \cap e : e \in \mathcal{E}\}|$, and by the Perles-Sauer-Shelah lemma, if H has VC-dimension d then $\pi_H(m) \leq \Sigma_{i=0}^{d} \binom{m}{i} = O(m^d)$. The dual hypergraph of a hypergraph $H = (V, \mathcal{E})$ is $H^* = (V^*, \mathcal{E}^*)$, where $V^* = \mathcal{E}$ and each $v \in V$ induces the hyperedge $e_v \in \mathcal{E}^*$, where $e_v = \{e \in \mathcal{E} : v \in e\}$. When the VC-dimension of H is d, the VC-dimension of H^* is denoted by d^*.

Any bipartite graph $G = G_{A,B}$ with vertex set $V(G) = A \cup B$ and edge set $E(G) \subset A \times B$, defines two hypergraphs: the primal hypergraph $H_G = (A, \mathcal{E}_B)$, where $\mathcal{E}_B = \{N(b) : b \in B\}$ is the collection of the open neighborhoods of the vertices in B, and the dual hypergraph $H_G^* = (B, \mathcal{E}_A)$, defined similarly. The VC-dimension of G is defined as the VC-dimension of H_G, and the dual VC-dimension of G is defined as the VC-dimension of H_G^*. The shatter function π_G and the dual shatter function π_G^* of G, are the shatter functions of H_G and of H_G^*, respectively.

In a remarkable result, Fox, Pach, Sheffer, Suk and Zahl [18] improved the bound of the Kővári-Sós-Turán theorem for graphs with VC-dimension at most d (for $d < t$). They showed:

▶ **Theorem 2** ([18]). *Let $t \geq 2$ and let $G_{A,B}$ be a bipartite graph with $|A| = m$ and $|B| = n$, satisfying $\pi_G(\ell) = O(\ell^d)$ and $\pi_G^*(\ell) = O(\ell^{d^*})$ for all ℓ. If G is $K_{t,t}$-free, then*

$$|E(G)| = O_{t,d,d^*}(\min\{mn^{1-\frac{1}{d}} + n, nm^{1-\frac{1}{d^*}} + m\}).$$

Theorem 2 spawned several follow-up papers. Janzer and Pohoata [22] obtained an improved bound of $o(n^{2-\frac{1}{d}})$ for graphs with VC-dimension d, where $m = n$ and $t \geq d > 2$, using the hypergraph removal lemma [20]. Do [14] and Frankl and Kupavskii [19] obtained improved bounds when t tends to infinity with n.

Improved bounds for Zarankiewicz's problem for incidence graphs. An *incidence graph* is a bipartite graph whose vertex set is a union of a set of points and a set of geometric objects, where the edges connect points to objects to which they are incident. Problems on incidence graphs are central in computational and combinatorial geometry. For example, the classical Erdős' unit distances problem asks for an upper bound on the number of edges in an incidence graph of points and unit circles. Furthermore, they are closely related to algorithmic problems in computational geometry, such as range searching and Hopcroft's problem (see [2, 12]).

Incidence graphs are naturally related to Zarankiewicz's problem. Indeed, the incidence graph of points and lines is $K_{2,2}$-free, and thus, the Kővári-Sós-Turán theorem implies that its number of edges is $O(n^{3/2})$. While this bound is tight for the incidence graph of the finite projective plane, the classical Szemerédi-Trotter theorem asserts that a stronger bound of $O(n^{4/3})$ holds in the real plane.

Motivated by this relation, Basit, Chernikov, Starchenko, Tao, and Tran [7] studied incidence graphs of points and axis-parallel boxes in $\mathbb{R}^d$, under the additional assumption that they are $K_{t,t}$-free. They obtained an $O(n \log^{2d} n)$ bound in $\mathbb{R}^d$, and a sharp $O(n\frac{\log n}{\log\log n})$ bound for dyadic axis-parallel rectangles in the plane. Independently, Tomon and Zakharov [31] obtained a weaker bound of $O(n \log^{2d+3} n)$ in $\mathbb{R}^d$ and a stronger bound of $O(n \log n)$ in the special case of a $K_{2,2}$-free incidence graph of points and axis-parallel rectangles in the plane.

At SODA'23, Chan and Har-Peled [11] initiated a systematic study of Zarankiewicz's problem for incidence graphs of points and various geometric objects. They obtained an $O(n(\frac{\log n}{\log\log n})^{d-1})$ bound for the incidence graph of points and axis-parallel boxes in $\mathbb{R}^d$ and observed that a matching lower bound construction appears in a classical paper of Chazelle ([13]; see also [30]). They also obtained an $O(n \log\log n)$ bound for points and pseudo-discs in the plane, and bounds for points and halfspaces, balls, shapes with "low union complexity", and more. The proofs in [11] use a variety of techniques, including shallow cuttings, a geometric divide-and-conquer, and biclique covers.

ϵ-nets and ϵ-t-nets. Given a hypergraph $H = (V, \mathcal{E})$ and $\epsilon > 0$, an *ϵ-net* for H is a set $S \subset V$ such that any hyperedge $e \in \mathcal{E}$ of size $\geq \epsilon|V|$ contains a vertex from S. The notion of ϵ-nets was introduced by Haussler and Welzl [21] who proved that any finite hypergraph with VC-dimension d admits an ϵ-net of size $O((d/\epsilon)\log(d/\epsilon))$ (a bound that was later improved to $O((d/\epsilon)\log(1/\epsilon))$ in [24]). ϵ-nets were studied extensively and have found applications in diverse areas of computer science, including machine learning, algorithms, computational geometry, and social choice (see, e.g., [3, 6, 8, 10]).

Very recently, Alon et al. [5] introduced the following notion of *ϵ-t-nets*, generalizing ϵ-nets and the notion of *ϵ-Mnets* that was studied by Mustafa and Ray [27] and by Dutta et al. [15]:

▶ **Definition 3.** *Let $\epsilon \in (0,1)$ and $t \in \mathbb{N} \setminus \{0\}$ be fixed parameters, and let $H = (V, \mathcal{E})$ be a hypergraph on n vertices. A set $N \subset \binom{V}{t}$ of t-tuples of vertices is called an ϵ-t-*net *if any hyperedge $e \in \mathcal{E}$ with $|e| \geq \epsilon n$ contains at least one of the t-tuples in N.*

Alon et al. [5] proved the following:

▶ **Theorem 4.** *For every $\epsilon \in (0,1)$, $C > 0$, and $t, d, d^* \in \mathbb{N} \setminus \{0\}$, there exists $C_1 = C_1(C, d^*)$ such that the following holds. Let H be a hypergraph on at least $C_1((t-1)/\epsilon)^{d^*}$ vertices with VC-dimension d and dual shatter function $\pi^*_H(m) \leq C \cdot m^{d^*}$. Then H admits an ϵ-t-net of size $O((d(1 + \log t)/\epsilon)\log(1/\epsilon))$, all of which elements are pairwise disjoint.*

In addition, the paper [5] studied the existence of small-sized ϵ-t-nets in various geometric settings in which it is known that the classical ϵ-net has size $O(1/\epsilon)$. In particular, they showed that the intersection hypergraph of two families of pseudo-discs and the dual incidence hypergraph of points and a family of regions with a linear union complexity, admit ϵ-2-nets of size $O(1/\epsilon)$, provided that they have at least $2/\epsilon$ vertices.

1.1 Our results

Zarankiewicz's problem via ϵ-t-nets. The basic observation underlying our results is a surprising connection between ϵ-t-nets and Zarankiewicz's problem. Consider a bipartite graph $G = G_{A,B} = (A \cup B, E)$ where $|A| = m, |B| = n$, and assume that for some ϵ, the primal hypergraph $H_G = (A, \mathcal{E}_B)$ (recall that $\mathcal{E}_B = \{N(b) : b \in B\}$ is the collection of the open neighborhoods of the vertices in B) admits an ϵ-t-net S of size s. We can partition the vertex set B into the set B' of "heavy" vertices that have degree at least ϵm in G, and the set B'' of "light" vertices that have degree less than ϵm.

We observe that the neighborhood $N(b)$ of any "heavy" vertex b must contain a t-tuple from S. Since G is $K_{t,t}$-free, any t-tuple from S is contained in at most $t-1$ neighborhoods $N(b_i)$. Hence, the number of "heavy" vertices is at most $s(t-1)$. This immediately yields the bound

$$|E(G)| \leq n\lfloor \epsilon m \rfloor + s(t-1)m,$$

where the first term is the contribution of the "light" vertices and the second term is the contribution of the "heavy" vertices.

Building upon that and several more observations, we develop a *recursive approach* which allows obtaining bounds for Zarankiewicz's problem using results on ϵ-t-nets (see Theorem 8 below).

A short proof of Fox et al.'s bound. Our first application of the ϵ-t-net approach is a short proof of the bound of Fox, Pach, Sheffer, Suk and Zahl [18] (Theorem 2 above). By combining the strategy described above with Theorem 4, we prove:

▶ **Theorem 5.** *Let $t \geq 2$ and let $G_{A,B}$ be a bipartite graph with $|A| = m$ and $|B| = n$, satisfying $\pi_G(\ell) = O(\ell^d)$ and $\pi^*_G(\ell) = O(\ell^{d^*})$ for all ℓ. If G is $K_{t,t}$-free, then we have*

$$|E(G)| = O_{t,d,d^*}(\min\{mn^{1-\frac{1}{d}} + n^{1+\frac{1}{d}}\log n, nm^{1-\frac{1}{d^*}} + m^{1+\frac{1}{d^*}}\log m\}).$$

The bound of Theorem 5 matches the bound of Theorem 2, whenever $d, d^* > 2$ and n, m do not differ "too much". Interestingly, when $n = m$ and $d^* = d$, our proof strategy can be combined with the proof strategy of Fox et al. [18] to obtain a slightly better bound. We omit the details due to space limitations.

Zarankiewicz's problem for intersection graphs. The *intersection graph* of a family $\mathcal{F}$ of geometric objects is a graph whose vertex set is $\mathcal{F}$, and whose edges connect pairs of objects whose intersection is non-empty. In the general (i.e., non-bipartite) setting, K_t-free intersection graphs of geometric objects were studied extensively, and have applications to the study of quasi-planar topological graphs (see, e.g., [17]).

Generalizing the systematic study of Zarankiewicz's problem for incidence graphs initiated by Chan and Har-Peled [11], we study Zarankiewicz's problem for *bipartite* intersection graphs of geometric objects – i.e., the maximum number of edges in a $K_{t,t}$-free graph $G_{A,B} = (A \cup B, E)$, where A, B are families of geometric objects, and objects $x \in A, y \in B$ are connected by an edge if their intersection is non-empty. Obviously, incidence graphs are the special case where A consists of a set of points.

It is important to note that this setting (i.e., bipartite intersection graphs) is different from the (standard) intersection graph of the family $A \cup B$, in which intersections inside A and inside B are also taken into account. The stark difference is exemplified well in the case of families A, B of segments in the plane. If the bipartite intersection graph of A, B

is $K_{2,2}$-free, the results of Fox et al. [18] imply the upper bound $O(n^{4/3})$ on its number of edges, and this bound is tight, in view of the tightness examples of the Szemerédi-Trotter theorem. On the other hand, if the (standard) intersection graph of $A \cup B$ is $K_{2,2}$-free, then the results of Fox and Pach [16] imply an upper bound of $O(n)$ on its number of edges (see also [26]). Our setting is the natural generalization of incidence graphs, in which only point-object incidences are taken into account, but not intersections between the objects.

A sharp bound for Zarankiewicz's problem for intersection graphs of pseudo-discs. A *family of pseudo-discs* is a family of simple closed Jordan regions in the plane such that the boundaries of any two regions intersect in at most two points. For example, a family of homothets (scaled translation copies) of a given convex body in the plane is a family of pseudo-discs. As a second application of the strategy described above, we obtain a linear upper bound for Zarankiewicz's problem for the intersection graph of two families of pseudo-discs.

▶ **Theorem 6.** *Let $t \geq 2$ and let $G = G_{A,B}$ be the bipartite intersection graph of families A, B of pseudo-discs, with $|A| = |B| = n$. If G is $K_{t,t}$-free then $|E(G)| = O(t^6 n)$.*

In fact, we show that the assertion of Theorem 6 holds (with a slightly weaker bound of $O(t^8 n)$) for a wider class of bipartite intersection graphs of any two families of so-called *non-piercing regions* – namely, families $\mathcal{F}$ of regions in the plane such that for any $S, T \in \mathcal{F}$, the region $S \setminus T$ is connected.

Theorem 6 improves and generalizes a result of Chan and Har-Peled [11] who obtained an upper bound of $O(n \log\log n)$ for the incidence graph of points and pseudo-discs. In order to prove Theorem 6 we show that the primal and the dual hypergraphs of G admit ϵ-t-nets of size $O_t(1/\epsilon)$ for all $n \geq \frac{2t}{\epsilon}$ (see Theorem 10). Thus, we also extend the results of [5], and believe that this might be of independent interest.

Theorem 6 demonstrates the added value of the new ϵ-t-net approach over previous approaches that used shallow cuttings. There are settings, like intersection graphs of pseudo-discs, for which one can show the existence of a linear-sized ϵ-t-net, while the existence of shallow cuttings is not known. In such settings, the ϵ-t-net approach yields stronger bounds than previous techniques.

An interesting problem which is left open is whether the dependence on t in Theorem 6 can be improved. It seems that the right dependence should be linear, like in the bounds of Chan and Har-Peled [11].

A sharp bound for intersection graphs of axis-parallel rectangles. The main class of incidence graphs studied in the previous papers [7, 11, 31] is incidence graphs of points and axis-parallel rectangles (and more generally, axis-parallel boxes in $\mathbb{R}^d$). We generalize this direction of study to the intersection graph of two families of axis-parallel rectangles, and obtain the following:

▶ **Theorem 7.** *Let $t \geq 2$, let $n \geq n_0$ for some $n_0(t)$, and let $G = G_{A,B}$ be the bipartite intersection graph of families A, B of axis-parallel rectangles in general position*[1]*, with $|A| = |B| = n$. If G is $K_{t,t}$-free, then $|E(G)| = O(tn \frac{\log n}{\log\log n})$.*

As follows from a lower bound given in [7], this result is sharp even in the special case where one of the families consists of points and the other consists of dyadic axis-parallel rectangles.

[1] The general position means that no two edges of rectangles in $A \cup B$ lie on the same vertical or horizontal line.

In the case of bipartite intersections graphs of families of axis-parallel rectangles, and even in the more basic case of incidence graphs of points and axis-parallel rectangles, the currently known bounds on the size of ϵ-t-nets do not allow obtaining efficient bounds for Zarankiewicz's problem using our ϵ-t-net based strategy. Indeed, among these settings, "small"-sized ϵ-t-nets for all $\epsilon \geq c/n$ are known to exist only for the incidence hypergraph of points and axis-parallel rectangles, and the size of the ϵ-t-net is $O(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\log\log\frac{1}{\epsilon})$ (see [5, Theorem 6.10]). Applying the strategy described above with an ϵ-t-net of such size would lead to an upper bound on the number of edges in a $K_{t,t}$-free incidence graph of points and axis-parallel rectangles, that is no better than $O(n\log n\log\log n)$.

In order to obtain the stronger (and tight) bound of $O(n\frac{\log n}{\log\log n})$ in the more general setting of bipartite intersection graphs of two families of axis-parallel rectangles, we combine the result of Chan and Har-Peled [11] with a combinatorial argument and with planarity arguments. Our technique allows us also to obtain a sharp $O(n)$ upper bound on the number of edges in the bipartite intersection graph of two families of n axis-parallel *frames* (i.e., boundaries of rectangles) in the plane, and an improved bound of $O(t^4 n)$ on the number of edges in the intersection graph of points and pseudo-discs (for which Chan and Har-Peled [11] obtained the bound $O(n\log\log n)$).

Organization of the paper. In Section 2 we present our new approach to Zarankiewicz's problem via ϵ-t-nets and prove Theorems 5 and 6. In Section 3 we obtain a sharp bound on the number of edges in a $K_{t,t}$-free bipartite intersection graph of two families of axis-parallel rectangles.

2 From ϵ-t-Nets to $K_{t,t}$-free Bipartite Graphs

2.1 A recursive upper bound on the size of $K_{t,t}$-free bipartite graphs

Let $G_{A,B}$ be a bipartite graph, where $|A| = m, |B| = n$. As was shown in the introduction, if for some $\epsilon > 0$, the primal hypergraph $H_G = (A, \{e_b = \{N(b)\}_{b\in B}\})$ admits an ϵ-t-net of size s, we may partition the vertex set B into the set B' of "heavy" vertices that have degree at least ϵm and the set B'' of "light" vertices that have degree less than ϵm, and observe that since $G_{A,B}$ is $K_{t,t}$-free, $|B'| \leq s(t-1)$. This yields the bound

$$|E(G)| \leq n\lfloor \epsilon m\rfloor + s(t-1)m,$$

where the first term is the contribution of the "light" vertices and the second term is the contribution of the "heavy" vertices.

If, in addition, the dual hypergraph $H_G^* = H(B, \{e_a = \{N(a)\}_{a\in A}\})$ admits an ϵ'-t-net of size s', then we can repeat the procedure described above and partition the vertex set A into the set A' of "heavy" vertices that have degree at least $\epsilon' n$ and the set A'' of "light" vertices that have degree less than $\epsilon' n$. By the same argument as above, $|A'| \leq s'(t-1)$. Hence, we can obtain the bound

$$|E(G)| \leq n\lfloor \epsilon m\rfloor + m\lfloor \epsilon' n\rfloor + ss'(t-1)^2,$$

where the first term bounds the contribution of the "light" vertices from B, the second term bounds the contribution of the "light" vertices from A, and the third term bounds the contribution of the "heavy"-"heavy" edges.

Furthermore, we can reduce the third term by observing that the set of "heavy"-"heavy" edges is the edge set of the bipartite graph $G_{A',B'}$ (i.e., the induced subgraph of G on the vertex set $A' \cup B'$), which is $K_{t,t}$-free, and thus, we may be able to apply to it the above partioning once again, provided that the corresponding primal and dual hypergraphs admit "small"-sized ϵ-t-nets. This gives rise to the recursive Algorithm 1 depicted below.

Algorithm 1 NumEdges.

Input: $G_{A,B}$, t
Output: Upper bound on $|E(G_{A,B})|$, assuming $G_{A,B}$ is $K_{t,t}$-free

1: Choose ϵ, ϵ'
2: Define s to be the minimum size of an ϵ-t-net for H_G
3: Define s' to be the minimum size of an ϵ'-t-net for H^*_G
4: Let $A' = \{v \in A : \deg_{G_{A,B}}(v) \geq \epsilon' n\}$
5: Let $B' = \{w \in B : \deg_{G_{A,B}}(w) \geq \epsilon m\}$
6: Return $n\lfloor \epsilon m \rfloor + m\lfloor \epsilon' n \rfloor$ +NumEdges($G_{A',B'}$)

Note that the choice of ϵ, ϵ' at Step 1 of the algorithm is not specified. The optimal choice is determined by the dependence of the size of the smallest ϵ-t-net of H_G and of H^*_G on ϵ. In the applications presented below, we simply choose both ϵ and ϵ' to be the smallest possible value for which the existence of a "small"-sized ϵ-t-net for the corresponding hypergraph is known.

Correctness of Algorithm 1. Let us call a vertex of $A' \cup B'$ *heavy*, and call the other vertices *light*. Algorithm 1 counts separately the edges of $G_{A,B}$ that involve a light vertex, and the edges that connect two heavy vertices. All the latter edges are counted by the recursion at Step 6.

Regarding the edges of $G_{A,B}$ that involve a light vertex, there are at most m light vertices in A, and each of them is involved in at most $\lfloor \epsilon' n \rfloor$ edges of $G_{A,B}$. Similarly, there are at most n light vertices in B, and each of them is involved in at most $\lfloor \epsilon m \rfloor$ edges of $G_{A,B}$. This explains the additive term $n\lfloor \epsilon m \rfloor + m\lfloor \epsilon' n \rfloor$ at Step 6.

Upper bound for Zarankiewicz's problem for hereditary classes of objects. Algorithm 1 allows establishing a recursive formula that yields an upper bound for Zarankiewicz's problem for a wide class of graphs. To present the formula in its full generality, a few more definitions and notations are needed.

For a bipartite graph $G = G_{A,B}$ where $|A| = m, |B| = n$, we denote by $f_G(m,k)$ the minimum size of a $\frac{k}{m}$-t-net of the primal hypergraph H that corresponds to G, and by $f^*_G(n,\ell)$ the minimum size of an $\frac{\ell}{n}$-t-net of the dual hypergraph H^* that corresponds to G.

We say that a class $\mathcal{F}$ of objects is *hereditary* if it is downwards closed, meaning that $(A \in \mathcal{F}) \wedge (A' \subset A) \Rightarrow (A' \in \mathcal{F})$. For example, the class of all families of pseudo-discs in the plane is clearly hereditary.

For two fixed hereditary classes of objects $\mathcal{F}, \mathcal{F}'$, we denote by $f(m,k) = f_{\mathcal{F},\mathcal{F}'}(m,k)$ and $f^*(n,\ell) = f^*_{\mathcal{F},\mathcal{F}'}(n,\ell)$ the maxima of $f_G(m,k)$ and of $f^*_G(n,\ell)$ (respectively) over all bipartite graphs $G = G_{A,B}$ such that $A \in \mathcal{F}$, $B \in \mathcal{F}'$, $|A| = m$, and $|B| = n$.[2] Furthermore, we denote by $g(m,n) = g_{\mathcal{F},\mathcal{F}'}(m,n)$ the maximum number of edges in a $K_{t,t}$-free bipartite graph $G_{A,B}$, where $A \in \mathcal{F}$, $B \in \mathcal{F}'$, $|A| = m$, and $|B| = n$.

[2] We note that an extra condition that $G_{A,B}$ is $K_{t,t}$-free could be added here, since all bipartite graphs encountered during the recursive process are $K_{t,t}$-free. It will be interesting to understand whether this additional assumption implies the existence of ϵ-t-nets of a smaller size.

▶ **Theorem 8.** *Let $\mathcal{F}, \mathcal{F}'$ be hereditary classes of objects. In the above notations, we have*

$$g(m,n) \leq \min_{1\leq k\leq m-1} \min_{1\leq \ell\leq n-1} ((k-1)n + (\ell-1)m + g((t-1)f(m,k),(t-1)f^*(n,\ell))). \quad (1)$$

Proof. For any bipartite graph $G_{A,B}$, where $A \in \mathcal{F}$, $B \in \mathcal{F}'$, $|A| = m$, and $|B| = n$, and any k, ℓ, we may apply Algorithm 1 with $\epsilon = \frac{k}{m}$ and $\epsilon' = \frac{\ell}{n}$, to obtain the bound

$$|E(G_{A,B})| \leq (k-1)n + (\ell-1)m + |E(G_{A',B'})|.$$

Here, the term $(k-1)n$ bounds the contribution of the "light" vertices in B, as there are at most n such vertices and each of them has degree strictly less than $\frac{k}{m}m = k$. The term $(\ell-1)m$ bounds the contribution of the "light" vertices in A in a similar way.

Note that we have $|B'| \leq (t-1)f(m,k)$. Indeed, let S be an ϵ-t-net for H_G of size $f(m,k)$. On the one hand, for each $b \in B'$, the hyperedge e_b contains a t-tuple from S (since $|e_b| \geq k$ and S is a $\frac{k}{m}$-t-net). On the other hand, as $G_{A,B}$ is $K_{t,t}$-free, any t-tuple in S participates in at most $t-1$ hyperedges of H_G. Thus, $|B'| \leq (t-1)|S| = (t-1)f(m,k)$. By the same argument, we have $|A'| \leq (t-1)f^*(n,\ell)$. Since $\mathcal{F}, \mathcal{F}'$ are hereditary, it follows that $|E(G_{A',B'})| \leq g((t-1)f(m,k),(t-1)f^*(n,\ell))$, and thus,

$$|E(G_{A,B})| \leq (k-1)n + (\ell-1)m + g((t-1)f(m,k),(t-1)f^*(n,\ell)).$$

Taking the minimum over all $1 \leq k \leq m-1$ and $1 \leq \ell \leq n-1$ and then the maximum over all $G_{A,B}$, where $A \in \mathcal{F}$, $B \in \mathcal{F}'$, $|A| = m$, and $|B| = n$, completes the proof. ◀

Theorem 8 allows leveraging results on ϵ-t-nets into upper bounds for Zarankiewicz's problem in a black box manner. The results presented in the following subsections are obtained by applying this approach (or parts of it) for specific classes of graphs.

2.2 Graphs with bounded VC-dimension

The first step of the approach presented above along with Theorem 4 yield a strikingly simple proof of Theorem 5.

Proof of Theorem 5. Put $\epsilon = \frac{C_1^{1/d^*}(t-1)}{m^{1/d^*}}$, where C_1 is the constant from Theorem 4. Let N be an ϵ-t-net for H_G of size $O((d(1+\log t)/\epsilon)\log(1/\epsilon))$, whose existence follows from Theorem 4.

Let $B' \subset B$ be the set of vertices with degree at least $\epsilon m = \Theta_{d^*,t}(m^{1-\frac{1}{d^*}})$ in G. We claim that

$$|B'| \leq (t-1)|N| = O_{d,t}(\frac{1}{\epsilon}\log\frac{1}{\epsilon}) = O_{d,d^*,t}(m^{\frac{1}{d^*}}\log m).$$

Indeed, on the one hand, for each $b \in B'$, the hyperedge e_b contains a t-tuple from N. On the other hand, as $G_{A,B}$ is $K_{t,t}$-free, any t-tuple in N participates in at most $t-1$ hyperedges of H_G. Thus, $|B'| \leq (t-1)|N|$, as asserted.

To complete the proof, we note that $|E(G)| = (\sum_{b\in B} d(b))$, where $d(b)$ is the degree of b in G. Hence, we have

$$|E(G)| = \sum_{b\in B} d(b) = \sum_{b\in B'} d(b) + \sum_{b\in B\setminus B'} d(b) \leq |B'|m + |B\setminus B'|\epsilon m = O_{d^*,d,t}(m^{1+1/d^*}\log m + nm^{1-1/d^*}).$$

The *min* assertion is achieved by applying the same argument to H_G^* instead of H_G. ◀

2.3 $K_{t,t}$-free bipartite intersection graphs of pseudo-discs

Our recursive strategy can be exploited optimally when both the primal and the dual hypergraphs that correspond to G admit $O_t(1/\epsilon)$-sized ϵ-t-nets for ϵ as small as $O(1/|V|)$, where $|V|$ is the number of vertices of the corresponding hypergraph. In this subsection we prove that this is the case for intersection graphs of two families of pseudo-discs.[3] Then, we use this to obtain an improved linear-sized bound on the number of edges of the graph.

We begin with a formal definition.

▶ **Definition 9.** *A family $\mathcal{F}$ of simple closed Jordan regions in $\mathbb{R}^2$ is called a* family of pseudo-discs *if for any $a, b \in \mathcal{F}$, the boundaries of a and b intersect at most twice.*

We prove the following ϵ-t-net theorem for a hypergraph induced by two families of pseudo-discs, which might be of independent interest. The theorem is "optimal", in the sense that it provides an $O_t(1/\epsilon)$-sized ϵ-t-net already when ϵn is constant.

▶ **Theorem 10.** *Let $\mathcal{F}_1$ and $\mathcal{F}_2$ be two families of pseudo-discs. Let H be a hypergraph whose vertex-set is $\mathcal{F}_1$, where each $b \in \mathcal{F}_2$ defines a hyperedge $e_b = \{a \in \mathcal{F}_1 : a \cap b \neq \emptyset\}$. If $|\mathcal{F}_1| = n$ and $\epsilon n \geq 2t$ then H admits an ϵ-t-net of size $O(t^5 \cdot \frac{1}{\epsilon})$.*

The proof of Theorem 10 makes use of the following definition and result:

▶ **Definition 11.** *Let $H = (V, \mathcal{E})$ be a hypergraph. The* Delaunay graph *of H is the graph on the same vertex-set, whose edges are the hyperedges of H of cardinality 2.*

▶ **Theorem 12** ([1], Theorem 6(ii,iii))**.** *Let $H = (V, \mathcal{E})$ be a hypergraph. Suppose there exists $C > 0$ such that for every $V' \subset V$, the Delaunay graph of the hypergraph induced by V' has less than $C|V'|$ hyperedges.*[4] *Denote the VC-dimension of H by d. Then:*

1. *$d \leq 2C$.*
2. *H has $O(t^{d-1}|V|)$ hyperedges of size at most t.*

Proof of Theorem 10. First, we find a "small" set $S \subset V(H) = \mathcal{F}_1$ such that each hyperedge in H of size at least ϵn contains at least t elements of S. To this end, we first let $K_1 \subset V(H)$ be an ϵ-net for H of size $O(\frac{1}{\epsilon})$. The existence of such a linear-sized ϵ-net is well-known (even for the more general case of two families of non-piercing regions), see, e.g., [5, Thm. 6.2]. Then, for each $2 \leq i \leq t$ sequentially, we consider the hypergraph induced on $V(H) \setminus (K_1 \cup \ldots \cup K_{i-1})$, and let K_i be an $O(\frac{1}{\epsilon})$-sized $\frac{\epsilon}{2}$-net for it. Let $S = \bigcup_{i=1}^{t} K_i$. Clearly, $|S| = O(\frac{t}{\epsilon})$.

We claim that each hyperedge $e \in \mathcal{E}(H)$ with $|e| \geq \epsilon n$ contains at least t elements of S. Indeed, if $|S \cap e| < t$, then there exists some $1 \leq i \leq t$ such that before the i'th step, the hyperedge induced by e contained less than t elements of $K_1 \cup \ldots \cup K_{i-1}$, and $e \cap K_i = \emptyset$. Since $\epsilon n \geq 2t$, we have $\epsilon n - t \geq \frac{\epsilon n}{2}$, a contradiction to K_i being an $\frac{\epsilon}{2}$-net.

Having the set $S \subset V(H)$ in hands, we construct an ϵ-t-net for H that consists of t-tuples in $\binom{S}{t}$. Let H_S be the hypergraph induced by H on S. It is known that the Delaunay graph of H, and therefore, of H_S, is planar – namely, the condition of Theorem 12 holds with $C = 3$, and the VC-dimension of H_S is at most 4 (see [23]). These two properties hold for any induced subhypergraph of H_S as well. By Theorem 12, this implies that any induced subhypergraph of H_S on m vertices contains $O(t^3 m)$ hyperedges of size at most t. By a simple double-counting argument, it follows that in any such induced subhypergraph, there exists a vertex that participates in $O(t^4)$ hyperedges of size at most t.

---

[3] The special case $t = 2$ of this result appeared in [5], but the proof method there is very specific for $t = 2$.

[4] The hypergraph induced by V' is $(V', \mathcal{E}')$, where $\mathcal{E}' = \{e \cap V' : e \in \mathcal{E}\}$.

Now we are ready to construct the desired ϵ-t-net, N. We choose a vertex a that participates in $O(t^4)$ hyperedges of size at most t of H_S, and add to N all the t-sized hyperedges (if exist) that contain a. Then we delete a, and continue inductively, in the same manner, with the hypergraph induced on $V(H_S) \setminus \{a\}$. We continue in this fashion until all vertices of S are removed. The number of steps is $|S| = O(\frac{t}{\epsilon})$, and each step contributes $O(t^4)$ t-tuples to N. Hence, in total, we have $|N| = O(t^5 \cdot \frac{1}{\epsilon})$.

We claim that N is an ϵ-t-net for H. Indeed, consider a hyperedge $e \in \mathcal{E}(H)$ with $|e| \geq \epsilon n$. By the construction of S, e contains at least t vertices from S. During the process in which we removed one-by-one the vertices of S, consider the step in which the size $|e \cap S|$ was reduced from t to $t-1$. At this step, e contained exactly t elements from S, that formed a t-tuple added to N. Hence, e contains a t-tuple from N. Thus, N is an ϵ-t-net for H. This completes the proof of Theorem 10. ◀

▶ Remark 13. It is clear from the proof of Theorem 10 that the theorem holds (up to a factor of $O_t(1)$) for any hypergraph H that satisfies the following properties:

1. Any induced subhypergraph $H' \subset H$ admits an ϵ-net of size $O(1/\epsilon)$, for any $\epsilon \geq \frac{t}{|V(H')|}$.
2. H has a hereditarily linear Delaunay graph.

An example of such a setting is the intersection hypergraph of two families of non-piercing regions in the plane.

The existence of a "small"-sized ϵ-t-net for the intersection hypergraph of pseudo-discs enables us to apply Algorithm 1, and thus to obtain an improvement and a generalization of the result of Chan and Har-Peled [11] mentioned in the introduction. The following lemma quantifies the partition of the vertices in steps 4-5 of Alg. 1 into "heavy" and "light" ones.

▶ **Lemma 14.** *Let $\mathcal{F}_1, \mathcal{F}_2$ be two families of pseudo-discs with $|\mathcal{F}_1| = |\mathcal{F}_2| = n$, and let $G = G_{\mathcal{F}_1,\mathcal{F}_2}$ be the bipartite intersection graph of $\mathcal{F}_1, \mathcal{F}_2$. If G is $K_{t,t}$-free and $\ell \geq 2t$, then the number of vertices in $\mathcal{F}_1 \cup \mathcal{F}_2$ whose degree in G is at least ℓ is $O(t^6 \frac{n}{\ell})$.*

Proof of Lemma 14. We prove the lemma w.l.o.g. for the vertices in $\mathcal{F}_2$. Let $\epsilon = \frac{\ell}{n}$. Since $\epsilon n = \ell \geq 2t$, we can apply Theorem 10 to obtain an ϵ-t-net N of size $O(\frac{t^5}{\epsilon})$ for the primal hypergraph H_G. Each hyperedge of H_G of size at least $\epsilon n = \ell$ contains a t-tuple from N, but since G is $K_{t,t}$-free, each such a t-tuple participates in at most $t-1$ hyperedges.

Therefore, the total number of hyperedges of size at least ℓ in $H_G = (\mathcal{F}_1, \mathcal{E}_{\mathcal{F}_2})$ is at most $(t-1)O(\frac{t^5}{\epsilon}) = O(\frac{t^6}{\epsilon}) = O(\frac{t^6 n}{\ell})$. This is exactly the number of vertices in $\mathcal{F}_2$ with degree at least ℓ in G. This completes the proof of Lemma 14. ◀

Now we are ready to prove Theorem 6. The idea of the proof is to apply Algorithm 1 with such a choice of ϵ, ϵ' in step 1 (which is actually done by choosing the parameter ℓ in Lemma 14), that the number of "heavy" vertices[5] in each part of the graph is reduced by a factor of 2. Intuitively, this factor-2-reduction saves the order of magnitude of the total sum from being affected by the number of steps in the recursive process. This choice can be made since (unlike in the general case of Theorem 4), Theorem 10 holds already when $\epsilon n \geq 2t$.

Proof of Theorem 6. Denote by $f(n)$ the maximum possible number of edges in $G_{\mathcal{F}_1,\mathcal{F}_2}$, for $\mathcal{F}_1, \mathcal{F}_2$ as in the statement of the theorem. Let $C \geq 1$ be a universal constant such that Lemma 14 holds with $C\frac{t^6 n}{\ell}$. We prove by induction that our claim holds with $f(n) \leq 8Ct^6 n$.

[5] The notions of "heavy" and "light" vertices were explained in the correctness proof of Algorithm 1.

For $n \leq 8Ct^6$, the assertion is trivial since $|E(G_{\mathcal{F}_1,\mathcal{F}_2})| \leq n^2$. We assume correctness for $\frac{n}{2}$ and prove the assertion for n. By Lemma 14 with $\ell = 2Ct^6 \geq 2t$, the number of vertices in $G_{\mathcal{F}_1,\mathcal{F}_2}$ with degree at least $2Ct^6$ is at most $C\frac{t^6 n}{\ell} = \frac{n}{2}$.

Recall that a vertex in $\mathcal{F}_1 \cup \mathcal{F}_2$ is called "heavy" if its degree in $G_{\mathcal{F}_1,\mathcal{F}_2}$ is at least ℓ, and otherwise, it is called "light". There are at most $n\ell$ edges in $G_{\mathcal{F}_1,\mathcal{F}_2}$ that connect a light vertex of $\mathcal{F}_1$ (resp., $\mathcal{F}_2$) with some vertex of $\mathcal{F}_2$ (resp., $\mathcal{F}_1$). The number of edges in $G_{\mathcal{F}_1,\mathcal{F}_2}$ that connect two heavy vertices is at most $f(\frac{n}{2})$, and by the induction hypothesis, $f(\frac{n}{2}) \leq 4Ct^6 n$. Therefore, the total number of edges in $G_{\mathcal{F}_1,\mathcal{F}_2}$ is at most $(2Ct^6+2Ct^6+4Ct^6)n = 8Ct^6 n$. ◀

▶ Remark 15. Theorem 6 can be readily generalized (albeit, with a slightly weaker bound of $O(t^8 n)$) to the more general setting of bipartite intersection graphs of two families of *non-piercing regions in the plane* – i.e., families $\mathcal{F}$ of regions such that for any $S, T \in \mathcal{F}$, $S \setminus T$ is connected. We omit the details due to space limitations.

3 $K_{t,t}$-free Bipartite Intersection Graphs of Axis-parallel Rectangles

In this section we prove Theorem 7 – a sharp upper bound on the number of edges in a $K_{t,t}$-free bipartite intersection graph of two families of axis-parallel rectangles. As was explained in the introduction, the current knowledge on ϵ-t-nets for intersection hypergraphs of families of axis-parallel rectangles is not sufficient for obtaining a sharp bound for Zarankiewicz's problem using our ϵ-t-net approach. Hence, we prove the theorem by an entirely different method that uses the sharp bound of Chan and Har-Peled on the number of edges in a $K_{t,t}$-free incidence graph of points and axis-parallel rectangles [11], along with other combinatorial and geometric techniques.

Let us restate the theorem, in a slighly different (but clearly equivalent) form.

▶ **Theorem 7.** *Let $t \geq 2$ and let $n, m \geq n_0$ for some $n_0(t)$. Let A, B be two families of axis-parallel rectangles, $|A| = n, |B| = m$, s.t. $A \cup B$ is in general position. If $G_{A,B}$ is $K_{t,t}$-free, then $|E(G_{A,B})| = O(t(n+m)\frac{\log(n+m)}{\log\log(n+m)})$.*

Proof. Any intersection between $a \in A$ and $b \in B$ belongs to exactly one of four types:

1. The rectangle a is strictly contained in the rectangle b.
2. The rectangle b is strictly contained in the rectangle a.
3. A vertical edge of b intersects a horizontal edge of a.
4. A vertical edge of a intersects a horizontal edge of b.

We bound separately the numbers of intersections of each of these types.

Intersections of type 1. We define a bipartite graph G whose vertices are all the corners of rectangles in A, and all the rectangles in B. A corner x is adjacent to a rectangle $b \in B$ if $x \in b$. Clearly, $|V(G)| = 4n + m$. We observe that G is $K_{4t-3,4t-3}$-free, since if some $4t-3$ corners are all contained in the same $4t-3$ rectangles of B, then these $4t-3$ corners belong to at least t different rectangles in A, and this contradicts the assumption that $G_{A,B}$ is $K_{t,t}$-free. Therefore, we can apply the following result of Chan and Har-Peled [11]:

▶ **Lemma 16** ([11], Lemma 4.4)**.** *Let P be a set of n points in $\mathbb{R}^2$, and let $\mathcal{R}$ be a family of m axis parallel rectangles in $\mathbb{R}^2$. If the incidence graph $G_{P,\mathcal{R}}$ is $K_{t,t}$-free, then $E(G_{P,\mathcal{R}}) = O_\epsilon(tn\frac{\log n}{\log\log n} + tm\log^\epsilon n)$, for any constant $\epsilon > 0$.*

Since the two sides of $V(G)$ contain $4n$ and m vertices, respectively, Lemma 16 (applied with any fixed $0 < \epsilon < 1$) yields

$$|E(G)| = O(t(n+m)\frac{\log(n+m)}{\log\log(n+m)}).$$

Note that each intersection of type 1 contributes exactly 4 edges to G. Hence, the number of intersections of type 1 is $O(t(n+m)\frac{\log(n+m)}{\log\log(n+m)})$. By symmetry, the same bound holds for the number of intersections of type 2.

Intersections of type 3. Define a bipartite graph $K = K_{S,S'}$ whose vertices are the horizontal edges of rectangles in A (that we call *horizontal vertices of* K), and the vertical edges of rectangles in B (that we call *vertical vertices of* K). A vertical vertex is adjacent to a horizontal vertex if the corresponding edges cross. Each intersection of type 3 contributes either 1, 2, or 4 edges to H. Therefore, the number of such intersections is at most $4|E(K)|$.

Denote the vertices of S (i.e., the vertical vertices of K) by $v_1, v_2, \ldots, v_m$ (in an arbitrary order). For $1 \le i \le 2m$, let d_i be the degree of v_i in K. Clearly, $|E(K)| = \Sigma_{i=1}^{2m} d_i$. Let $\mathcal{F} \subset \binom{S'}{2t-1}$ be the family of all *canonical* $(2t-1)$*-tuples* of horizontal vertices, where a $(2t-1)$-tuple T of horizontal vertices is called canonical if there exists some vertical segment L (not necessarily a vertical vertex!) that intersects exactly the vertices of T among all the horizontal vertices (i.e., we have $\{x \in S' : x \cap L \neq \emptyset\} = T$).

▷ Claim 17. In the above notations, $|\mathcal{F}| = O(t^5 n)$.

We leave the proof of Claim 17 to the end of this section, and continue with the proof of Theorem 7 (assuming the claim).

For each $1 \le i \le 2m$, we define x_i to be the number of canonical $(2t-1)$-tuples of horizontal vertices which the vertical vertex v_i intersects. That is,

$$x_i = |\{\{h_1, \ldots, h_{2t-1}\} \in \mathcal{F} : \forall 1 \le j \le 2t-1, v_i \cap h_j \neq \emptyset\}|.$$

We would like to obtain lower and upper bounds on $\Sigma_{i=1}^{2m} x_i$.

On the one hand, $\Sigma_{i=1}^{2m} x_i \le (2t-2)|\mathcal{F}|$. Indeed, for any canonical $(2t-1)$-tuple $\{h_1, \ldots, h_{2t-1}\} \in \mathcal{F}$, at most $2t-2$ v_i's intersect all of $h_1, \ldots, h_{2t-1}$, since otherwise, $G_{A,B}$ contains $K_{t,t}$ as a subgraph (as the h_j's must belong to at least t different rectangles in A and the v_i's must belong to at least t different rectangles in B). By Claim 17, this implies

$$\Sigma_{i=1}^{2m} x_i = (2t-2) \cdot O(t^5 n) = O(t^6 n). \tag{2}$$

On the other hand, for each $1 \le i \le 2m$, we have $d_i - 2t + 2 \le x_i$. Indeed, if $d_i \le 2t-2$ the inequality is trivial. If $d_i \ge 2t-1$ and v_i intersects (w.l.o.g.) the horizontal vertices $h_1, \ldots, h_{d_i}$ in this order, then each consecutive $(2t-1)$-subsequence of $h_1, \ldots, h_{d_i}$ belongs to $\mathcal{F}$, since it is the intersection of some subsegment of v_i with the set of horizontal vertices. Therefore,

$$\Sigma_{i=1}^{2m}(d_i - 2t + 2) \le \Sigma_{i=1}^{2m} x_i. \tag{3}$$

Combining (2) and (3) together, we obtain

$$\Sigma_{i=1}^{2m}(d_i - 2t + 2) = O(t^6 n),$$

and thus,

$$|E(K)| = \Sigma_{i=1}^{2m} d_i = O(t^6 n + tm).$$

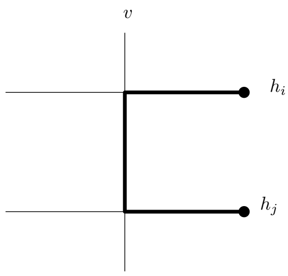

Figure 1 The planar drawing of $v = \{h_i, h_j\}$ (in bold).

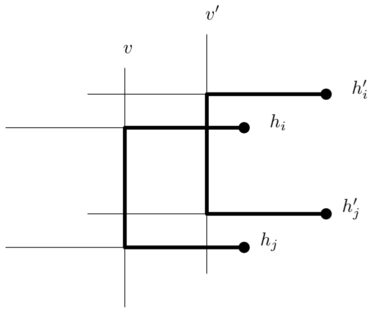

Figure 2 In this figure, the planar drawing of $\{h_i, h_j\}$ intersects the planar drawing of $\{h'_i, h'_j\}$. However, v' intersects three horizontal segments, and therefore $\{h'_i, h'_j\} \notin E(Del(J))$.

Hence, the number of intersections of type 3 is $O(t^6 n + tm)$, which is negligible compared to $O(t(n+m)\frac{\log(n+m)}{\log\log(n+m)})$ for any fixed t. By symmetry, the same bound applies to the number of intersections of type 4. This completes the proof of the theorem (assuming Claim 17). ◀

The only part that remains is the proof of Claim 17.

Proof of Claim 17. Define a hypergraph J whose vertices are the horizontal edges of the rectangles in $A - h_1, \ldots, h_{2n}$, and each vertical segment v (which is not necessarily a vertical vertex!) defines a hyperedge e_v which is the subset of $\{h_1, \ldots, h_{2n}\}$ that v intersects.

Note that $\mathcal{F}$ is the set of hyperedges of size $2t-1$ of J. We would like to bound the size of this set by $O(t^5 n)$. To this end, we prove that the Delaunay graph of J, $Del(J)$, has a hereditarily linear number of edges. It is clearly sufficient to prove that $|E(Del(J))|$ is linear in $|V(Del(J))| = 2n$.

Let us describe a planar drawing of $Del(J)$. We represent each vertex of $Del(J)$ by the right endpoint of the corresponding horizontal edge. Each edge $v = \{h_i, h_j\} \in E(Del(J))$ is drawn as a 3-polygonal path that starts at the right endpoint of h_i, continues on the subsegment of v that connects h_i and h_j, and continues on h_j towards its right endpoint (see Figure 1).

The drawing described here is a planar drawing of $Del(J)$. It is easy to verify that if for some four distinct vertices h_i, h_j, h'_i, h'_j, the planar drawing of $v = \{h_i, h_j\}$ intersects the planar drawing of $v' = \{h'_i, h'_j\}$, then either v or v' cannot be an edge of $Del(J)$ (see Figure 2). As $Del(J)$ admits a planar drawing in which no two vertex-disjoint edges intersect, it is planar (e.g., by an easy special case of the Hanani-Tutte theorem).

Thus, each subgraph of $Del(J)$ on ℓ vertices contains at most $3\ell - 6$ edges. By Theorem 12, this implies that the number of hyperedges of size at most $2t-1$ in J is $O(t^5 n)$. In particular, $|\mathcal{F}| = O(t^5 n)$, as asserted. ◁

▶ Remark 18. Note that some intersections between two rectangles $a \in A, b \in B$ were counted twice. In the counting of type 1, we actually counted all the intersections in which a corner of a is contained in b, and not only the ones in which $a \subseteq b$.

▶ Remark 19. The proof of Theorem 7 readily implies that the number of edges in a $K_{t,t}$-free bipartite intersection graph $G_{A,B}$ of two families of axis-parallel frames (i.e., boundaries of rectangles) with $|A| = n, |B| = m$ is $O(t^6 n + tm)$. Indeed, the only possible types of intersections between a pair of axis-parallel frames are intersections of types 3, 4 presented above. In the proof, the number of such intersections is bounded by $O(t^6 n + tm)$.

References

1. E. Ackerman, B. Keszegh, and D. Pálvölgyi. On tangencies among planar curves with an application to coloring L-shapes, proceedings of EuroComb 2021, Trends in Mathematics — Research Perspectives, CRM Barcelona 14, 2021.
2. P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. In *Advances in Discrete and Computational Geometry. Ed. by B. Chazelle, J. E. Goodman, and R. Pollack. AMS Press*, pages 1–56, 1999.
3. N. Alon, G. R. Brightwell, H. A. Kierstead, A. V. Kostochka, and P. Winkler. Dominating sets in k-majority tournaments. *J. Comb. Theory, Ser. B*, 96(3):374–387, 2006.
4. N. Alon, D. Haussler, and E. Welzl. Partitioning and geometric embedding of range spaces of finite vapnik-chervonenkis dimension. In D. Soule, editor, *Proceedings of SoCG'1987*, pages 331–340. ACM, 1987.
5. N. Alon, B. Jartoux, C. Keller, S. Smorodinsky, and Y. Yuditsky. The ϵ-t-net problem. *Discret. Comput. Geom.*, 68(2):618–644, 2022. doi:10.1007/s00454-022-00376-x.
6. S. Arya, G. Dias da Fonseca, and D. M. Mount. Approximate polytope membership queries. *SIAM J. Comput.*, 47(1):1–51, 2018.
7. A. Basit, A. Chernikov, S. Starchenko, T. Tao, and C.-M. Tran. Zarankiewicz's problem for semilinear hypergraphs. *Forum Math. Sigma*, 9:Paper No. e59, 23, 2021.
8. A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Learnability and the vapnik-chervonenkis dimension. *J. ACM*, 36(4):929–965, 1989.
9. W. G. Brown. On graphs that do not contain a Thomsen graph. *Canadian Math. Bulletin*, 9(3):281–285, 1966.
10. T. M. Chan. Improved deterministic algorithms for linear programming in low dimensions. *ACM Trans. Algorithms*, 14(3):30:1–30:10, 2018.
11. T. M. Chan and S. Har-Peled. On the number of incidences when avoiding an induced biclique in geometric settings. In *Proceedings of SODA 2023*, pages 1398–1413. SIAM, 2023. doi:10.1137/1.9781611977554.ch50.
12. T. M. Chan and D. W. Zheng. Hopcroft's problem, log-star shaving, 2d fractional cascading, and decision trees. In *Proceedings of SODA 2022*, pages 190–210. SIAM, 2022.
13. B. Chazelle. Lower bounds for orthogonal range searching: I. The reporting case. *J. ACM*, 37(2):200–212, 1990.
14. T. Do. Representation complexities of semi-algebraic graphs. *SIAM J. Discret. Math.*, 4(33):1864–1877, 2019.
15. K. Dutta, A. Ghosh, B. Jartoux, and N. H. Mustafa. Shallow packings, semialgebraic set systems, Macbeath regions, and polynomial partitioning. *Discret. Comput. Geom.*, 61(4):756–777, 2019.
16. J. Fox and J. Pach. Separator theorems and Turán-type results for planar intersection graph. *Advances in Mathematics*, 219(3):1070–1080, 2008.
17. J. Fox and J. Pach. Coloring $\mathrm{k_k}$-free intersection graphs of geometric objects in the plane. *Eur. J. Comb.*, 33(5):853–866, 2012.
18. J. Fox, J. Pach, A. Sheffer, A. Suk, and J. Zahl. A semi-algebraic version of Zarankiewicz's problem. *J. Euro. Math. Soc.*, 19(6):1785–1810, 2017.

19 N. Frankl and A. Kupavskii. On the Erdős-Purdy problem and the Zarankiewitz problem for semialgebraic graphs, available at arxiv:2112.10245, 2021.

20 W. T. Gowers. Hypergraph regularity and the multidimensional szemerédi's theorem. *Annals of Mathematics*, 166:897–946, 2007.

21 D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.

22 O. Janzer and C. Pohoata. On the Zarankiewicz problem for graphs with bounded VC-dimension, available at arxiv: 2009.00130, 2020.

23 B. Keszegh. Coloring intersection hypergraphs of pseudo-disks. In Bettina Speckmann and Csaba D. Tóth, editors, *Proceedings of SoCG 2018*, volume 99 of *LIPIcs*, pages 52:1–52:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.SoCG.2018.52.

24 J. Komlós, J. Pach, and G. J. Woeginger. Almost tight bounds for epsilon-nets. *Discret. Comput. Geom.*, 7:163–173, 1992.

25 P. Kővári, V. Sós, and P. Turán. On a problem of Zarankiewicz. *Colloq. Math.*, 3:50–57, 1954.

26 N. H. Mustafa and J. Pach. On the Zarankiewicz problem for intersection hypergraphs. *J. Comb. Theory, Ser. A*, 141:1–7, 2016.

27 N. H. Mustafa and S. Ray. ϵ-Mnets: Hitting geometric set systems with subsets. *Discret. Comput. Geom.*, 57(3):625–640, 2017.

28 N. H. Mustafa and K. Varadarajan. *Epsilon-approximations & epsilon-nets, in: Handbook of Discrete and Computational Geometry, 3rd ed*, pages 1241–1267. CRC Press, Boka Raton, 2018.

29 B. Sudakov. Recent developments in extremal combinatorics: Ramsey and Turán type problems. In *Proceedings of the International Congress of Mathematicians. Volume IV*, pages 2579–2606, 2010.

30 I. Tomon. Coloring lines and delaunay graphs with respect to boxes. *Random Structures & Algorithms*, 2023.

31 I. Tomon and D. Zakharov. Turán-type results for intersection graphs of boxes. *Comb. Probab. Comput.*, 30(6):982–987, 2021.