

Matchings vs hitting sets among half-spaces in low dimensional euclidean spaces

Shay Moran* Rom Pinchasi†

July 9, 2015

Abstract

Let \mathcal{F} be any collection of linearly separable sets of a set P of n points either in \mathbb{R}^2 , or in \mathbb{R}^3 . We show that for every natural number k either one can find k pairwise disjoint sets in \mathcal{F} , or there are $O(k)$ points in P that together hit all sets in \mathcal{F} . The proof is based on showing a similar result for families \mathcal{F} of sets separable by pseudo-discs in \mathbb{R}^2 . We complement these statements by showing that analogous result fails to hold for collections of linearly separable sets in \mathbb{R}^4 and higher dimensional euclidean spaces.

*Department of Computer Science, Technion—Israel Institute of Technology, Haifa 32000, Israel, and Max Planck Institute for Informatics, Saarbrücken, Germany. shaymoran1@gmail.com.

†Mathematics Department, Technion—Israel Institute of Technology, Haifa 32000, Israel. room@math.technion.ac.il. Supported by ISF grant (grant No. 1357/12).

1 Introduction

Let $\mathcal{H} = (V, E)$ be a hyper-graph. A hitting set for \mathcal{H} is a subset of vertices which intersects every edge in E . A matching in \mathcal{H} is a subset of mutually disjoint edges. Let $\tau(\mathcal{H})$ denote the size of a minimum hitting set of \mathcal{H} and let $\nu(\mathcal{H})$ denote the size of a maximum matching of \mathcal{H} . The parameters $\tau(\mathcal{H}), \nu(\mathcal{H})$ were studied extensively in combinatorics and in computer science. $\tau(\mathcal{H})$ and $\nu(\mathcal{H})$ relate to each other. Indeed, every hitting set must contain a distinct element from each edge in any matching and therefore $\nu(\mathcal{H}) \leq \tau(\mathcal{H})$. Moreover, by strong duality for linear programming it follows that the size of a minimum *fractional*¹ hitting set, denoted by $\tau^*(\mathcal{H})$, is equal to the size of a maximum *fractional*² matching, denoted by $\nu^*(\mathcal{H})$. So every hyper-graph \mathcal{H} satisfies:

$$\nu(\mathcal{H}) \leq \nu^*(\mathcal{H}) = \tau^*(\mathcal{H}) \leq \tau(\mathcal{H}).$$

Hyper-graphs \mathcal{H} for which $\tau(\mathcal{H}) = \nu(\mathcal{H})$ or for which $\tau(\mathcal{H})$ and $\nu(\mathcal{H})$ are close to each other have also been studied. See for example [4, 2, 3] and references within.

We study the gap between $\nu(\mathcal{H})$ and $\tau(\mathcal{H})$ for hyper-graphs \mathcal{H} which can be realized by an arrangement of half-spaces in \mathbb{R}^d when d is small. This property is quantified by the *affine sign-rank*. The affine sign-rank of a hyper-graph \mathcal{H} is the minimum number d for which there is an identification of $V(\mathcal{H})$ as points in \mathbb{R}^d and of $E(\mathcal{H})$ as half-spaces in \mathbb{R}^d such that for all $v \in V(\mathcal{H}), e \in E(\mathcal{H}), v \in e$ if and only if the point corresponding to v is in the half-space corresponding to e . The affine sign-rank is closely related³ to the sign-rank of \mathcal{H} which was studied in many contexts such as geometry [5], machine learning [14, 7, 15], communication complexity [23, 12, 13, 26] and more.

Hyper-graphs with small affine sign-rank have small VC dimension (at most the affine sign-rank plus one) and therefore, by [8, 11], for such hyper-graphs:

$$\tau(\mathcal{H}) \leq O(\tau^*(\mathcal{H}) \log \tau^*(\mathcal{H})).$$

How about $\nu(\mathcal{H})$? Is it also close to $\nu^*(\mathcal{H})$? In general, low VC dimension does not imply that $\nu(\mathcal{H})$ is close to $\nu^*(\mathcal{H})$. A simple example is given by $\mathcal{H} = (P, L)$ where P and L are the sets of points and lines in a projective plane of order n . Recall that in a projective plane of order n $|P| = |L| = n^2 + n + 1$, each two lines intersect in a unique point, each two points have a unique line containing both of them, each line contain exactly $n + 1$ points and each point has exactly $n + 1$ lines containing it. Thus, its VC dimension is 2, $\nu(\mathcal{H}) = 1$ (since every two lines intersect) and $\nu^*(\mathcal{H}) \geq \frac{|L|}{n+1} = \frac{n^2+n+1}{n+1} = \Omega(n)$ as we may choose a $\frac{1}{n+1}$ fraction of every line so that every point is covered exactly once and the total weight of the fractional matching is $\frac{|L|}{n+1}$. However, since the affine sign-rank of \mathcal{H} is $\Omega(n^{1/2})$ [13, 5] this example does not rule out the possibility that τ and ν are close for hyper-graphs of constant affine sign-rank.

We show that if the affine sign-rank of \mathcal{H} is less than 4 then $\tau(\mathcal{H}) = \Theta(\nu(\mathcal{H}))$. We complement this by showing that there are hyper-graphs \mathcal{H} with affine sign-rank 4 such that $\nu(\mathcal{H}) = 1$ and $\tau(\mathcal{H})$ is arbitrarily large.

We note that the fact that $\tau(\mathcal{H}) = \Theta(\nu(\mathcal{H}))$ when the affine sign-rank is 2 is already known [10]. For completeness we add our alternative proof for it and show how this proof is generalized to capture the case of affine sign-rank 3.

¹put a non-negative weight on each vertex so that for every edge, the total weight of all vertices in it is at least 1

²put a non-negative weight on each edge so that for every vertex, the total weight of all edges covering it is at most 1

³The affine sign-rank is between the sign-rank and the sign-rank plus 1.

2 Our results

For a set P of points in \mathbb{R}^d and a family \mathcal{F} of ranges in \mathbb{R}^d we denote by $\mathcal{H}(P, \mathcal{F})$ the hyper-graph on the set of vertices P whose edges consist of the sets $\{P \cap F \mid F \in \mathcal{F}\}$, without multiplicities. So, the affine sign-rank of \mathcal{H} is d if and only if there is a set P of points in \mathbb{R}^d and a family \mathcal{F} of half-spaces in \mathbb{R}^d such that \mathcal{H} is isomorphic to $\mathcal{H}(P, \mathcal{F})$.

2.1 The case of affine sign-rank 2 and pseudo-discs

As mentioned above, we show that if \mathcal{H} is a hyper-graph with affine sign-rank 2 then $\tau(\mathcal{H}) = \Theta(\nu(\mathcal{H}))$. In fact, we prove it for a more general class of hyper-graphs: A family \mathcal{C} of simple closed curves in \mathbb{R}^2 is called a family of pseudo-circles if every two curves in \mathcal{C} are either disjoint or cross at two points. A family of circles, no two of which touch, is a natural example for such a family. A family of pseudo-discs is a family of compact sets whose boundaries form a family of pseudo-circles. Natural examples for families of pseudo-discs are translates of a fixed convex set in the plane as well as homothetic copies of a fixed convex set in the plane.

Note that if the affine sign-rank of \mathcal{H} is 2 then there is a set of points P in the plane and a family of pseudo-discs \mathcal{F} such that \mathcal{H} is isomorphic $\mathcal{H}(P, \mathcal{F})$ (just replace each half-space by a large enough circular disc).

Theorem 1 ([10]). *Let P be a set of points in the plane and let \mathcal{F} be a family of pseudo-discs. Let \mathcal{H} be the hyper-graph $\mathcal{H} = \mathcal{H}(P, \mathcal{F})$. Then for every integer $k \geq 1$ either \mathcal{H} has k pairwise disjoint edges, or one can find $O(k)$ points in P that hit all the edges in \mathcal{H} .*

Theorem 1 implies that every \mathcal{H} with affine sign-rank 2 has $\tau(\mathcal{H}) = \Theta(\nu(\mathcal{H}))$. Theorem 1 was proved by Chan and Har-Peled in [10], however the proof that we present here is based on a different approach. Our methods are useful also in the case when the affine sign-rank is 3. The proof of Theorem 1 is based on the following Theorem:

Theorem 2. *Let \mathcal{F} be a family of pseudo-discs in the plane. Let P be a finite set of points in the plane and consider the hyper-graph $\mathcal{H} = \mathcal{H}(P, \mathcal{F})$. There exists an edge e in \mathcal{H} such that the maximum cardinality of a matching among the edges in \mathcal{H} that intersect with e is at most 156.*

Theorem 2 implies Theorem 1 as follows. Apply Theorem 2 to find an edge e in \mathcal{H} such that among those edges intersecting it there are at most 156 pairwise disjoint ones. Delete e and those edges intersecting it from \mathcal{H} . Repeat this until the graph is empty. If this continues k steps, then we find k pairwise disjoint edges. Otherwise, we decompose \mathcal{H} into less than k families, $\mathcal{H}_1, \dots, \mathcal{H}_\ell$, of edges such that in each family \mathcal{H}_i there are at most 156 pairwise disjoint edges.

We then show that for every $1 \leq i \leq \ell$ the edges in \mathcal{H}_i can be pierced by $O(1)$ points. This will conclude the proof of Theorem 1. In order to show that each \mathcal{H}_i is indeed pierced by $O(1)$ points, we rely on the techniques of Alon and Kleitman in [4] by proving a (p, q) Theorem for each of the \mathcal{H}_i (see the proof of Theorem 1).

Theorem 2 is a discrete version (and therefore also generalization) of Theorem 1 in [24], in which the set P is the entire plane. The proof of Theorem 2 follows the proof of Theorem 1 in [24] with some suitable adjustments.

The result in Theorem 2 (and also Theorem 1 in [24]) can be interpreted as saying that in every family of pseudo-discs there is a so called “small” pseudo-disc. Indeed, notice that in every

family of circular discs, the disc of smallest area, D , has the property that the maximum number of mutually disjoint discs from the family that intersect with it is at most $O(1)$ (see the introduction in [24] and the references therein for more details). Theorem 2 implies that the same phenomenon happens in every family of pseudo-discs.

The authors of [10], in which Theorem 1 was first proved, explicitly note that one of the challenges they overcome is the absence of a “smallest pseudo-disc”. In this paper and in [24] the existence of such pseudo-disc is proved. We prove Theorems 1 and 2 in Section 3.

2.2 The case of affine sign-rank 3

Theorem 3. *Let P be a set of points in \mathbb{R}^3 and let \mathcal{F} be a family of half-spaces. Let \mathcal{H} be the hyper-graph $\mathcal{H} = \mathcal{H}(P, \mathcal{F})$. Then for every integer $k \geq 1$ either \mathcal{H} has k pairwise disjoint edges, or one can find $O(k)$ points in P that hit all the edges in \mathcal{H} .*

Like in the case of affine sign-rank 2, the proof of Theorem 3 is based on the following theorem that is an analogue of Theorem 2:

Theorem 4. *Let P be a set of points in \mathbb{R}^3 and let \mathcal{F} be a family of half-spaces. Let \mathcal{H} be the hyper-graph $\mathcal{H} = \mathcal{H}(P, \mathcal{F})$. Then there exists an edge in \mathcal{H} such that the cardinality of the maximum matching among the edges in \mathcal{H} intersecting it is at most 156.*

We prove Theorems 3 and 4 in Section 4.

2.3 The case of affine sign-rank 4

We show that the analogous result to Theorems 1 and 3 fails for affine sign-rank greater than 3.

Theorem 5. *For every $n \in \mathbb{N}$ There exists a set P of $N = \binom{n}{2}$ points and a set \mathcal{F} of n half-spaces in \mathbb{R}^4 such that:*

1. *Every two edges in $\mathcal{H}(P, \mathcal{F})$ have a non-empty intersection (which implies that $\nu(H) = 1$).*
2. *Any subset of P which pierce all edges in $\mathcal{H}(P, \mathcal{F})$ has at least $\frac{n-1}{2}$ points in it (i.e. $\tau(H) \geq \frac{n-1}{2}$).*

We prove Theorem 5 in Section 5

2.4 Connection to ϵ -nets

Theorems 1 and 3 immediately imply a result from [21] about the existence of an ϵ -net of size linear in $\frac{1}{\epsilon}$ for hyper-graphs $\mathcal{H}(P, \mathcal{F})$, where \mathcal{F} is a family of pseudo-discs in \mathbb{R}^2 (hence also the special case where \mathcal{F} is a family of half-planes) or half-spaces in \mathbb{R}^3 . Indeed, given such a hyper-graph \mathcal{H} and $\epsilon > 0$, we delete from \mathcal{H} all the edges of cardinality smaller than $\epsilon|P|$. Set $k = \frac{1}{\epsilon}$. Notice that now \mathcal{H} does not contain k pairwise disjoint edges simply because every edge is of cardinality greater than $\epsilon|P|$. It follows that one can find $O(k) = O(\frac{1}{\epsilon})$ points in P that meet all the edges in \mathcal{H} .

Pach and Tardos [22] have recently shown that for every $\epsilon > 0$ and large enough n , there is a collection of n points, P , in \mathbb{R}^4 and a collection of half spaces, \mathcal{F} , such that every ϵ -net for $\mathcal{H}(P, \mathcal{F})$ has size $\Omega(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$. This corresponds to Theorem 5, and in fact implies some variant of it.

2.5 An algorithmic application

An immediate algorithmic application of Theorems 1 and 3 is a polynomial constant factor approximation algorithm for finding maximum matching in hyper-graphs of the form $\mathcal{H}(\mathcal{P}, \mathcal{F})$ where \mathcal{F} is a set of pseudo-discs (or half-planes) and $\mathcal{P} \subseteq \mathbb{R}^2$ or \mathcal{F} is a set of half-spaces in \mathbb{R}^3 and $\mathcal{P} \subseteq \mathbb{R}^3$. Indeed, given such a hyper-graph \mathcal{H} , we can repeatedly find a “small” edge $e \in E(\mathcal{H})$ in the sense of Theorems 1 and 3, add it to the matching and then delete e and those edges intersecting it from \mathcal{H} and continue until all the edges of \mathcal{H} are consumed. The final maximal (with respect to set containment) matching M has size which is at least $\frac{1}{156}$ of the size of a maximum matching. We note that Chan and Har-Peled [10] give a PTAS for maximum matching among pseudo-discs, with a different constant, also for the weighted case.

3 The case of affine sign-rank 2 and pseudo-discs

In this section we prove Theorem 1 and Theorem 2.

We start with the proof of Theorem 2 and then use this result to prove Theorem 1.

An important special case of Theorem 2 in which the set P is the set of all point in \mathbb{R}^2 is shown in [24]. The proof of Theorem 2 will follow the same lines of the proof in [24] with some suitable adjustments.

The idea of the proof is to show that if B is a maximum matching in \mathcal{H} then on average over all edges $e \in B$ the cardinality of a maximum matching among the edges in \mathcal{H} that intersects with e is less than 157. This means that there exists an edge in B with the desired property.

We will make use of the following lemma that is in fact Corollary 1 in [24]:

Lemma 1. *Let B be a family of pairwise disjoint sets in the plane and let \mathcal{F} be a family of pseudo-discs. Let D be a member of \mathcal{F} and suppose that D intersects exactly k members of B one of which is the set $e \in B$. Then for every $2 \leq \ell \leq k$ there exists a set $D' \subset D$ such that D' intersects e and exactly $\ell - 1$ other sets from B , and $\mathcal{F} \cup \{D'\}$ is again a family of pseudo-discs.*

We will also need the next lemma that is parallel to (and will take the place of) Lemma 2 in [24].

Lemma 2. *Let \mathcal{F} be a family of pseudo-discs in the plane. Let P be a finite set of points in the plane and consider the hyper-graph $\mathcal{H} = \mathcal{H}(P, \mathcal{F})$. Assume B is a subgraph of \mathcal{H} consisting of pairwise disjoint hyper-edges. Consider the graph G whose vertices correspond to the edges in B and connect two vertices $e, e' \in B$ by an edge if there is an edge in \mathcal{H} that has a nonempty intersection with e and with e' and has an empty intersection with all other edges in B . Then G is planar.*

Proof. We draw G as a topological graph in the plane as follows. From every edge $e \in B$ we pick one vertex, that we denote by $v(e)$, and the collection of all these vertices is the set V of vertices

of G . Denote by \mathcal{H}_2 the set of all edges in H that have a non-empty intersection with precisely two of the edges in B . For every pair of edges e and e' in B that are intersected by some edge f (possibly such an edge f is not unique) in \mathcal{H}_2 we draw an edge between $v(e)$ and $v(e')$ as follows. Pick a vertex $x \in e \cap f$ and a vertex $x' \in e' \cap f$. Recall that f is the intersection of P with some pseudo-disc D in \mathcal{F} . Similarly, let C and C' be two pseudo-discs in \mathcal{F} whose intersection with P is equal to e and e' , respectively. Let $W_{xx'}$ be an arc, connecting x and x' , that lies entirely in D . Let $W_{v(e)x}$ be an arc connecting $v(e)$ to x that lies entirely in C . Let $W_{v(e')x'}$ be an arc connecting $v(e')$ to x' that lies entirely in C' . Finally, we draw the edge in G connecting $v(e)$ and $v(e')$ as the union (or concatenation) of $W_{v(e)x}$, $W_{xx'}$, and $W_{x'v(e')}$. We will show that any two edges in G that do not share a common vertex are drawn so that they cross an even number of times. The Hanani-Tutte Theorem ([16, 28]) then implies the planarity of G .

We will use the following elementary lemma from [9]:

Lemma 3 (Lemma 1 in [9]). *Let D_1 and D_2 be two pseudo-discs in the plane. Let x and y be two points in $D_1 \setminus D_2$. Let a and b be two points in $D_2 \setminus D_1$. Let γ_{xy} be any Jordan arc connecting x and y that is fully contained in D_1 . Let γ_{ab} be any Jordan arc connecting a and b that is fully contained in D_2 . Then γ_{xy} and γ_{ab} cross an even number of times.*

Let $v(e), v(e')$ and $v(k), v(k')$ be four distinct vertices of G . This means in particular that e, e', k , and k' are four pairwise disjoint hyper-edges in B . Suppose that $v(e)$ and $v(e')$ are connected by an edge in G . This means that there are $x \in e$ and $x' \in e'$ and $f \in \mathcal{H}_2$ such that $x \in e \cap f$ and $x' \in e' \cap f$. Let E, E' , and F in \mathcal{F} be the pseudo-discs such that $e = E \cap P$, $e' = E' \cap P$, and $f = F \cap P$. Suppose also that $v(k)$ and $v(k')$ are connected by an edge in G . This means that there are $y \in k$ and $y' \in k'$ and $q \in \mathcal{H}_2$ such that $y \in k \cap q$ and $y' \in k' \cap q$. Let K, K' , and Q in \mathcal{F} be the pseudo-discs such that $k = K \cap P$, $k' = K' \cap P$, and $q = Q \cap P$.

By Lemma 3, $W_{v(e)x}$ and $W_{v(g)y}$ cross an even number of times. Indeed, E contains $v(e)$ and x and does not contain $v(k)$ and y . K contains $v(k)$ and y and does not contain $v(e)$ and x . Similarly, each of $W_{v(e)x}$, $W_{xx'}$, and $W_{v(e')x'}$ crosses each of $W_{v(k)y}$, $W_{yy'}$, and $W_{v(k')y'}$ an even number of times. We conclude that the edge in G connecting $v(e)$ and $v(e')$ crosses the edge in G connecting $v(k)$ and $v(k')$ an even number of times, as desired. ■

Proof of Theorem 2. The proof goes almost verbatim as the proof of Theorem 1 in [24]. Lemma 2 in [24] is replaced by the above Lemma 2.

Let B be a collection of pairwise disjoint edges in \mathcal{H} of maximum cardinality and let $n = |B|$. For every $e \in B$ denote by $\alpha_1(e)$ the size of a maximum matching among those edges in \mathcal{H} that intersect with e but with no other edge in B . Denote by $\alpha_2(e)$ the size of a maximum matching among those edges in \mathcal{H} that intersect with e and with precisely one more edge in B . Denote by $\alpha_3(e)$ the size of a maximum matching among those edges in \mathcal{H} that intersect with e and with at least two more edges in B . Observe that it is enough to show that $\sum_{e \in B} \alpha_1(e) + \alpha_2(e) + \alpha_3(e) < 157n$.

We first note that for every $e \in B$ we must have $\alpha_1(e) \leq 1$. Indeed, otherwise one can find two disjoint edges e' and e'' in H that do not intersect with any edge in B but e . The set $B \cup \{e', e''\} \setminus \{e\}$ contradicts that maximality of B .

Next, we show that $\sum_{e \in B} \alpha_2(e) \leq 12n$. Consider the graph G whose vertices correspond to the edges in B and connect two vertices $e, e' \in B$ by an edge if there is an edge in \mathcal{H} that has a nonempty intersection with e and with e' and has an empty intersection with all other edges in B . By Lemma 2, G is planar. Therefore, G has at most $3n$ edges. For every $e \in B$ denote by $d(e)$ the

degree of e in G . Therefore,

$$\sum_{e \in B} d(e) \leq 6n. \quad (1)$$

We claim that for every e in B we have $\alpha_2(e) \leq 2d(e)$. Indeed, otherwise by the pigeonhole principle one can find three pairwise disjoint edges g, g' , and g'' in \mathcal{H} and an edge e' in B such that each of g, g' , and g'' intersects e and e' but no other edge in B . In this case $B \cup \{g, g', g''\} \setminus \{e, e'\}$ contradicts the maximality of B .

Inequality (1) implies now $\sum_{e \in B} \alpha_2(e) \leq 12n$. It remains to show that $\sum_{e \in B} \alpha_3(e) < 144n$. The derivation of this inequality is more involved than the derivation of the inequalities regarding α_1, α_2 . We will show that if it is not the case that $\sum_{e \in B} \alpha_3(e) < 144n$, then we can derive an (impossible) embedding of $K_{3,3}$ in the plane.

Denote by \mathcal{F}_3 the subfamily of \mathcal{F} that consists of pseudo-discs in \mathcal{F} that intersect with three or more edges in B . Using repeatedly Lemma 1 with $\mathcal{F} = \mathcal{H}_3$ and with $\ell = 3$, we can find, for every $D \in \mathcal{H}_3$ and every $e \in B$ that is intersected by D , a (new) pseudo-disc $D^e \subset D$ that intersects with e and with exactly two more sets from B . Moreover, the collection of all the new sets D^e obtained in this way is a family of pseudo-discs. We denote this family of pseudo-discs by \mathcal{D} . Let T denote the set of all triples of edges in B that are intersected by a pseudo-disc in \mathcal{D} .

We denote by Z the collection of all pairs of sets from B that appear together in some triple in T . We claim that $|Z| < 12n$: Pick every set in B with probability $\frac{1}{2}$. Call a pair $\{e, e'\}$ in Z *good* if both e and e' were picked and an edge $f \in B$ such that e, e' , and f is a triple in T was not picked. The expected number of good pairs in Z is at least $1/8$ of the pairs in Z . On the other hand, by Lemma 2 the set of good pairs in Z is the set of edges of a planar graph (on an expected number of $n/2$ vertices) and therefore the expected number of good pairs is less than $3 \cdot \frac{n}{2}$.

Now consider the graph K whose set of vertices is the edges in B and whose set of edges is Z . For every $e \in B$ denote by $d(e)$ the degree of e in this graph. Notice that, in view of the above, $\sum_{e \in B} d(e) = 2|Z| < 24n$.

Fix $e \in B$. Define a graph K^e on the set of neighbors of e in K where we connect two neighbors e_1, e_2 of e in K by an edge in K^e if and only if $\{e, e_1, e_2\}$ is a triple in T . This is equivalent to that there is $D \in \mathcal{D}$ that intersects with e, e_1 , and with e_2 . Denote by $t(e)$ the number of edges in K^e . By ignoring the set e and applying Lemma 2, we see that K^e is planar. K^e has $d(e)$ vertices and is planar and therefore $t(e) < 3d(e)$.

We claim that for every $e \in B$ we must have $\alpha_3(e) \leq 2t(e)$. Indeed, assume to the contrary that $\alpha_3(e) > 2t(e)$. Then there is a collection Q of at least $2t(e) + 1$ pairwise disjoint edges of \mathcal{H} , each of which has a non-empty intersection with e and with at least two more edges in B . Because of Lemma 1, every edge in Q must have a non-empty intersection with e and with at least two edges e' and e'' that form a pair in Z . The hyper-edges e' and e'' are therefore connected by an edge in K^e . By the pigeonhole principle, because there are only $t(e)$ edges in K^e while $|Q| \geq 2t(e) + 1$, there exist e' and e'' that are connected by an edge in K^e such that e, e' , and e'' are all intersected by three (pairwise disjoint) edges $g_1, g_2, g_3 \in \mathcal{D}$. This is impossible as it gives an embedding of the graph $K_{3,3}$ in the plane. To see this, recall that also the sets e, e_1, e_2 are pairwise disjoint. For every $1 \leq i, j \leq 3$ add a small pseudo-disc surrounding one point in the intersection of e_i and g_j . Lemma 2 implies now an (impossible) embedding of $K_{3,3}$ in the plane.

We conclude that

$$\sum_{e \in \mathcal{B}} \alpha_3(e) < \sum_{e \in \mathcal{B}} 2t(e) \leq \sum_{e \in \mathcal{B}} 6d(e) \leq 6 \cdot 24n = 144n.$$

The proof is now complete as we have

$$\sum_{e \in \mathcal{B}} \alpha_1(e) + \alpha_2(e) + \alpha_3(e) < n + 12n + 144n = 157n$$

This implies the existence of $e \in \mathcal{B}$ such that $\alpha_1(e) + \alpha_2(e) + \alpha_3(e) \leq 156$. ■

Having proved Theorem 2, we are now ready to prove Theorem 1.

Proof of Theorem 1. Repeatedly apply Theorem 2 and find an edge e in \mathcal{H} such that among those edges intersecting it there are at most 156 pairwise disjoint ones. Then delete e and those edges intersecting it from \mathcal{H} and continue. If we can continue k steps, then we find k pairwise disjoint edges. Otherwise, we decompose \mathcal{H} into less than k families, $\mathcal{H}_1, \dots, \mathcal{H}_\ell$, of edges such that in each family \mathcal{H}_i there are at most 156 pairwise disjoint edges.

We will now show that for every $1 \leq i \leq \ell$ the edges in \mathcal{H}_i can be pierced by $O(1)$ points. This will conclude the proof of Theorem 1.

Our strategy is to show that the edges in \mathcal{H}_i have the so called (p, q) property for some p and q . That is, out of every p sets in \mathcal{H}_i there are q that have a non-empty intersection. In fact, by the definition of \mathcal{H}_i , it has the $(157, 2)$ property because there are at most 156 sets in \mathcal{H}_i that are pairwise disjoint. This is the first step. The next step is to show a (p, q) (for the same q above, that is $q = 2$) theorem for hyper-graphs $\mathcal{H}(\mathcal{P}, \mathcal{F})$ where \mathcal{F} is a family of pseudo-discs. This means that we will need to show that for a family of pseudo-discs \mathcal{F} if $\mathcal{H}(\mathcal{P}, \mathcal{F})$ has the (p, q) property, then one can find a constant number of points in \mathcal{P} that together pierce all edges in $\mathcal{H}(\mathcal{P}, \mathcal{F})$.

In order to complete the second step we will rely on the techniques of Alon and Kleitman in [4]. Rather than repeating their proof and adjusting it to our case, we observe, following Alon et. al in [3] and Matoušek in [20] that it is enough to show that the edges of $\mathcal{H}(\mathcal{P}, \mathcal{F})$ have fractional Helly number 2 (see below) and have a finite VC-dimension, which implies the existence of an ϵ -net of size that depends only on ϵ . These two ingredients are enough to show that $\mathcal{H}(\mathcal{P}, \mathcal{F})$ has a $(p, 2)$ theorem for every $p > 2$.

We recall that a hyper-graph \mathcal{H} is said to have a fractional Helly number k if for every $\alpha > 0$ there is $\beta > 0$ such that for any n and any collection of n sets in \mathcal{F} in which there are at least $\alpha \binom{n}{k}$ k -tuples that have nonempty intersection one can find a point incident to at least βn of the sets. Here β may depend only on α (and the hyper-graph \mathcal{H}) but not on n . In our setting the hyper-graph \mathcal{H} is of the form $\mathcal{H}(\mathcal{P}, \mathcal{F})$ where \mathcal{F} is a set of pseudo discs and \mathcal{P} is a set of points. We will see that every such \mathcal{H} has fractional Helly number 2 and that the corresponding β does not depend on \mathcal{P} nor on \mathcal{F} (it will only depend on certain combinatorial properties that are possessed by every family of pseudo-discs).

We recall also the notion of *union complexity* of a family of sets. We denote by $U_{\mathcal{F}}(m)$ the maximum complexity (that is, number of faces of all dimensions) of the boundary of the union of any m members of \mathcal{F} . We will need the following well known result from [18] saying that for a family \mathcal{F} of pseudo-discs we have $U_{\mathcal{F}}(m) \leq 12m$

We will use the following theorem from [25] (see Theorem 1 there) relating the notion of fractional Helly number with that of union complexity.

Theorem 6. *Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a function such that $\lim_{x \rightarrow \infty} g(x) = 0$. Suppose that \mathcal{F} is a family of geometric objects in \mathbb{R}^d in general position, (that is, no point belongs to the intersection of more than d boundaries of sets in \mathcal{F}) such that $U_{\mathcal{F}}(m) \leq g(m)m^k$ for every $m \in \mathbb{N}$. Then for every set of points P the family \mathcal{F}_P has fractional Helly number at most k and this is in a way that depends only on the function g and not on \mathcal{F} or P .*

To be more precise, for every $\alpha > 0$ there is a $\beta > 0$ such that for any family \mathcal{F} satisfying the conditions in the theorem and a set of points P in \mathbb{R}^d the following is true: For any collection of n sets in $\mathcal{H}(\mathcal{F}, P)$ in which there are at least $\alpha \binom{n}{k}$ k -tuples that have nonempty intersection one can find a point in P incident to at least βn of the sets.

Theorem 6 (with $d = 2$ and $k = 2$) and the linear bound on the union complexity of pseudo-discs [18] imply that $\mathcal{H}(P, \mathcal{F})$ has fractional Helly number at most 2. (Notice that we may assume without loss of generality that the sets in \mathcal{F} are indeed in general position and therefore Theorem 6 applies here.)

It is well known and not hard to show (see for example Theorem 9 in [9]) that for a family \mathcal{F} of pseudo-discs and a set P of points the hyper-graph $\mathcal{H}(P, \mathcal{F})$ has a bounded VC-dimension (in fact at most 3). Therefore, each \mathcal{H}_i has an ϵ -net of size that depends only on ϵ (see [17]). The method of Alon and Kleitman [4] implies that each $\mathcal{H}(P, \mathcal{F})$ satisfies a $(p, 2)$ theorem. That is, if any subset S of edges in $\mathcal{H}(P, \mathcal{F})$ satisfies the $(p, 2)$ property (from every p sets in S there are 2 sets that intersect), then there are $c(p)$ (a constant that depends only on p) vertices that together pierce all the sets in S (see Theorem 4 and the discussion around it in [20]).

By our assumption each, \mathcal{H}_i has the $(p, 2)$ -property for $p = 157$. It follows that one can find a set of points of cardinality at most $c(157)k$ that together intersect all the edges in \mathcal{H} . ■

4 The case of half-spaces in \mathbb{R}^3 .

In this section we prove Theorem 3. The proof follows the same trajectory as the proof of Theorem 1 with analogous lemmata. Technically, the challenge in this case is to derive the analogous lemmata for half-spaces in \mathbb{R}^3 .

For the proof of Theorem 3 we will need a corresponding three dimensional version of Lemma 2:

Lemma 4. *Let \mathcal{F} be a family of half-spaces in \mathbb{R}^3 . Let P be a finite set of points in \mathbb{R}^3 and consider the hyper-graph $\mathcal{H} = \mathcal{H}(P, \mathcal{F})$. Assume B is a subgraph of \mathcal{H} consisting of pairwise disjoint hyper-edges. Consider the graph G whose vertices correspond to the edges in B and connect two vertices $e, e' \in B$ by an edge if there is an edge in \mathcal{H} that has a nonempty intersection with e and with e' and has an empty intersection with all other edges in B . Then G is planar.*

Proof. We notice that if the points of P are in (strictly) convex position, then Lemma 4 follows almost right away from Lemma 2. To see this let S denote the convex hull of P and for every half-space F in \mathcal{F} let F^S denote the intersection of F with the boundary of S . Then the collection $\{F^S \mid F \in \mathcal{F}\}$ is a family of pseudo-discs lying on the boundary of S . Now Lemma 4 follows from Lemma 2 that, although stated in the plane, applies also to the boundary of S (homeomorphic to the two dimensional sphere).

When the points of P are not in convex position such a simple reduction is not possible anymore. Nevertheless, we will be able to make use of Lemma 2 after some suitable modifications.

Denote by M the union of all edges in B . We say that a point of M is *extreme* if it lies on the boundary of the convex hull of M .

Lemma 5. *Let e_1 and e_2 be two edges in B . Suppose that there exists an edge $f \in \mathcal{H}$ such that f has a nonempty intersection with e_1 and with e_2 and f does not intersect any other edge in B . Then there exists a half-space F' , not necessarily in \mathcal{F} , such that both intersections of F' with e_1 and with e_2 contain extreme points of M and still F' does not intersect any other edge in B but e_1 and e_2 .*

Proof. We shall use the following basic fact several times: Any half-space that has a non-empty intersection with M contains an extreme point of it. Let F denote the half-space in \mathcal{F} such that $f = F \cap P \supset F \cap M$. F contains at least one extreme vertex of M . Because $F \cap M \subset e_1 \cup e_2$ we conclude that there is an extreme vertex of M either in $F \cap e_1$, or in $F \cap e_2$ (if there is an extreme vertex of M in both, then we are done with $F' = F$). Without loss of generality assume that $F \cap e_2$ contains an extreme vertex of M . Let $E_1 \in \mathcal{F}$ be the half-space such that $e_1 = E_1 \cap P$. E_1 contains an extreme vertex of M that belongs to e_1 . Let ℓ denote the line of intersection of the boundaries of F and E_1 . Notice that $(F \cup E_1) \cap M \subset e_1 \cup e_2$. Take $F' = F$ and start rotating F' about the line ℓ such that at each moment $F' \subset F \cup E_1$. At each moment of the rotation until F' coincides with E_1 , the half-space F' contains the intersection $F \cap E_1$ and therefore F' has a nonempty intersection with e_1 . We stop at the last moment where F' still contains an extreme vertex of M that belongs to e_2 . At this moment F' must also contain a vertex of e_1 that is extreme in M . This is because at each moment F' must contain an extreme vertex of M . This completes the proof of the lemma. ■

Going back to the proof of Lemma 4, let S denote the convex hull of M . For every edge e in B let $F(e) \in \mathcal{F}$ be the half-space in \mathcal{F} such that $e = F(e) \cap P$. Denote by \tilde{e} the set of extreme vertices of M in e . Notice that for every $e \in B$ we have $\tilde{e} \neq \emptyset$ because every edge in B is the intersection of P with some half-space (in \mathcal{F}). Let \tilde{M} denote the set of extreme points in M . Because M is just the union of all edges in B , we have $\tilde{M} = \cup_{e \in B} \tilde{e}$. Observe that $\{\tilde{e} \mid e \in B\}$ is the set of edges of the hyper-graph $\tilde{\mathcal{H}} = \mathcal{H}(\tilde{M}, \{F(e) \mid e \in B\})$. For every pair of hyper edges $e, e' \in B$ that are neighbors in the graph G (defined in the statement of Lemma 4) let $F(e, e') \in \mathcal{F}$ denote some half-space in \mathcal{F} that has a nonempty intersection only with the edges e and e' from B . By Lemma 5, there exists a half-space that, with a slight abuse of notation, we denote by $F(\tilde{e}, \tilde{e}')$, not necessarily in \mathcal{F} , such that $F(\tilde{e}, \tilde{e}')$ has a non-empty intersection only with \tilde{e} and with \tilde{e}' from the collection $\{\tilde{f} \mid f \in B\}$.

Let

$$\mathcal{F}' = \{F_e \mid e \in B\} \cup \{F(\tilde{e}, \tilde{e}') \mid (e, e') \text{ is an edge in } G\}.$$

We define now a graph G' whose set of vertices is $B' = \{\tilde{e} \mid e \in B\}$. We connect \tilde{e} and \tilde{e}' in B' by an edge in G' if there is an edge f in the hyper-graph $\mathcal{H}(\tilde{M}, \mathcal{F}')$ such that f has a nonempty intersection with \tilde{e} and with \tilde{e}' and f has an empty intersection with all other sets in B' . It follows from the discussion above that if e and e' are two sets in B that are connected by an edge in G , then \tilde{e} and \tilde{e}' in B' are connected by an edge in G' .

Because \tilde{M} is in convex position, the hyper-graph $\mathcal{H}(\tilde{M}, \mathcal{F}')$ can be presented as a hyper-graph on the set of vertices \tilde{M} whose set of edges correspond to pseudo-discs on S , where S is the boundary of the convex hull of M . We then apply Lemma 2 (where B is replaced by $\{\tilde{e} \mid e \in B\}$ and \mathcal{F} is

replaced by \mathcal{F}') and conclude that G' is planar. The planarity of G follows because G is a subgraph of G' . ■

We are now ready to prove Theorem 4. The proof will follow the lines and will have a similar structure as of the proof of the corresponding theorem for pseudo-discs in the plane, namely Theorem 2.

Proof of Theorem 4. As in the proof of Theorem 2, let B be a maximum (in cardinality) collection of pairwise disjoint edges in \mathcal{H} and let $n = |B|$. For every $e \in B$ denote by $\alpha_1(e)$ the maximum cardinality of a matching among those edges in \mathcal{H} that intersect with e but with no other edge in B . Denote by $\alpha_2(e)$ the maximum cardinality of a matching among those edges in \mathcal{H} that intersect with e and with precisely one more edge in B . Denote by $\alpha_3(e)$ the maximum cardinality of a matching among those edges in \mathcal{H} that intersect with e and with at least two more edges in B . It is enough to show that $\sum_{e \in B} \alpha_1(e) + \alpha_2(e) + \alpha_3(e) < 157n$.

For every $e \in B$ we must have $\alpha_1(e) \leq 1$, or else we get a contradiction to the maximality of B (as in the proof of Theorem 2).

Next we show that $\sum_{e \in B} \alpha_2(e) \leq 12n$. Consider the graph G whose vertices correspond to the edges in B and connect two vertices $e, e' \in B$ by an edge if there is an edge in \mathcal{H} that has a nonempty intersection with e and with e' and has an empty intersection with all other edges in B . By Lemma 4, G is planar. Therefore, G has at most $3n$ edges. For every $e \in B$ denote by $d(e)$ the degree of e in G . Therefore,

$$\sum_{e \in B} d(e) \leq 6n. \quad (2)$$

We claim that for every e in B we have $\alpha_2(e) \leq 2d(e)$. Indeed, otherwise, by the pigeonhole principle, one can find three pairwise disjoint edges g, g' , and g'' in \mathcal{H} and an edge $e' \in B$ such that each of g, g' , and g'' intersects e and e' but no other edge in B . In this case $B \cup \{g, g', g''\} \setminus \{e, e'\}$ contradicts the maximality of B . Inequality (2) implies now $\sum_{e \in B} \alpha_2(e) \leq 12n$.

It remains to show that $\sum_{e \in B} \alpha_3(e) < 144n$. Denote by \mathcal{F}_3 the subfamily of \mathcal{F} that consists of half-spaces in \mathcal{F} that intersect with three or more edges in B . Like in the proof of Theorem 2 this part is more involved. Similarly, we will show that if it is not the case that $\sum_{e \in B} \alpha_3(e) < 144n$, then we derive an (impossible) embedding of $K_{3,3}$ in an arrangement of hyper-planes in \mathbb{R}^3 (see Claim 1).

For every $F \in \mathcal{F}_3$ and every $e \in B$ that is intersected by F , we find a (new) half-space F^e that intersects with e and with exactly two more edges in B . To do this, let $v \in e$ be an extreme vertex of P and let h be a hyper-plane supporting the convex hull of P at v . Let ℓ denote the line of intersection of h and the boundary of F . Rotate F about the line ℓ until F intersects only three edges in B one of which must be e because at all times of rotation we have $v \in F$.

We denote the family of all new half-spaces obtained this way by \mathcal{D} . Let T denote the set of all triples of edges in B that are intersected by half-spaces in \mathcal{D} .

We denote by Z the collection of all pairs of sets from B that appear together in some triple in T . One can show that $|Z| < 12n$: Pick every set in B with probability $\frac{1}{2}$. Call a pair $\{e, e'\}$ in Z *good* if both e and e' were picked and an edge $f \in B$ such that e, e' , and f is a triple in T was not picked. The expected number of good pairs in Z is at least $1/8$ of the pairs in Z . On the other hand, by Lemma 4 the set of good pairs in Z is the set of edges of a planar graph (on an expected

number of $n/2$ vertices). (We refer the reader to the proof of Theorem 2 to see this argument a bit more detailed.)

Now consider the graph K whose set of vertices is the edges in B and whose edges are those pairs in Z . For every $e \in B$ denote by $d(e)$ the degree of e in this graph. Notice that, in view of the above, $\sum_{e \in B} d(e) = 2|Z| < 24n$.

Fix $e \in B$. Define a graph K^e on the set of neighbors of e in K where we connect two neighbors e_1, e_2 of e in K by an edge in K^e if and only if $\{e, e_1, e_2\}$ is a triple in T . This is equivalent to that there is $D \in \mathcal{D}$ that intersects with e, e_1 , and with e_2 . Denote by $t(e)$ the number of edges in K^e . By ignoring the set e and applying Lemma 4, we see that K^e is planar. K^e has $d(e)$ vertices and is planar and therefore $t(e) < 3d(e)$.

We claim that for every $e \in B$ we must have $\alpha_3(e) \leq 2t(e)$.

Indeed, assume to the contrary that $\alpha_3(e) > 2t(e)$. Then there is a collection Q of at least $2t(e) + 1$ pairwise disjoint edges of H , each of which has a non-empty intersection with e and with at least two more edges in B . Every edge in Q has a non-empty intersection with e and with at least two edges e' and e'' that form a pair in Z . The hyper-edges e' and e'' are therefore connected by an edge in K^e . By the pigeonhole principle, because there are only $t(e)$ edges in K^e while $|Q| \geq 2t(e) + 1$, there exist e' and e'' that are connected by an edge in K^e such that e, e' , and e'' are all intersected by three (pairwise disjoint) edges $g_1, g_2, g_3 \in Q \subset \mathcal{H}$. We claim that this situation is impossible. This follows directly from the following claim

Claim 1. *It is impossible to find three half-spaces u_1, u_2, u_3 in \mathbb{R}^3 and another three half-spaces w_1, w_2, w_3 such that there are nine points q_{ij} for $1 \leq i, j \leq 3$ satisfying q_{ij} lies only in u_i and w_j from the half-spaces $u_1, u_2, u_3, w_1, w_2, w_3$.*

Proof. Considering the dual problem, it is enough to show that one cannot find three points u_1, u_2, u_3 in \mathbb{R}^3 and another three points $w_1, w_2, w_3 \in \mathbb{R}^3$ such that there for every $1 \leq i, j \leq 3$ there is a half-space containing only u_i and w_j from the points $u_1, u_2, u_3, w_1, w_2, w_3$.

Without loss of generality we assume that all the points are in general position. We may also assume that one of the triangles $\Delta u_1 u_2 u_3$ or $\Delta w_1 w_2 w_3$ is not a face of the convex hull of $\{u_1, u_2, u_3, w_1, w_2, w_3\}$. Otherwise, the points $u_1, u_2, u_3, w_1, w_2, w_3$ are in convex position and each of the segments $[u_i, w_j]$ is an edge of this convex polytope (because by assumption each pair of vertices w_i, u_j is separable from the rest of the vertices by a hyper-plane). The skeleton graph of a three dimensional convex polytope is planar and therefore cannot contain $K_{3,3}$ as a subgraph. Therefore, without loss of generality we assume that that the hyper-plane through u_1, u_2 , and u_3 separates two of the points w_1, w_2 , and w_3 . Let h denote this hyper-plane and assume without loss of generality that w_1 and w_2 lie above h while w_3 lies below h . We observe that the line through w_1 and w_2 must cross triangle $\Delta u_1 u_2 u_3$ for otherwise u_1, u_2, u_3, w_1, w_2 are in convex position and the edge-graph of their convex hull is the non-planar K_5 . Without loss of generality assume that w_1 lies closer than w_2 to triangle $\Delta u_1 u_2 u_3$. Denote by O the point of intersection of the line through w_1 and w_2 with h . For $i = 1, 2, 3$ let Q_i be a half-space containing only w_1 and u_i from $u_1, u_2, u_3, w_1, w_2, w_3$. Observe that all three half-spaces Q_1, Q_2 , and Q_3 must contain the point O (as they separate w_1 and w_2) and, assuming h is horizontal, their supporting hyper-planes must all lie above O . This implies that Q_1, Q_2 , and Q_3 cover the whole half-space below h which is impossible as none of Q_1, Q_2 , and Q_3 may contain w_3 . ■

Remark. Although it is tempting to believe that the collection of all 2-sets (that is, sets of two points separable by a half-space) of a set of points in \mathbb{R}^3 is the set of edges of a planar graph, this

is not the case. One can check that K_5 can be realized in this way. Claim 1 shows that $K_{3,3}$ cannot be realized in this way.

Going back to the proof of Theorem 4, we have:

$$\sum_{e \in \mathcal{B}} \alpha_3(e) < \sum_{e \in \mathcal{B}} 2t(e) \leq \sum_{e \in \mathcal{B}} 6d(e) \leq 6 \cdot 24n = 144n.$$

The proof is now complete as we have

$$\sum_{e \in \mathcal{B}} \alpha_1(e) + \alpha_2(e) + \alpha_3(e) < n + 12n + 144n = 157n,$$

and this implies the existence of $e \in \mathcal{B}$ such that $\alpha_1(e) + \alpha_2(e) + \alpha_3(e) \leq 156$. ■

In the same way that Theorem 1 is a corollary of Theorem 2, we conclude Theorem 3 from Theorem 4.

Proof of Theorem 3. Repeatedly apply Theorem 4 and find an edge e in \mathcal{H} such that among those edges intersecting it there are at most 156 pairwise disjoint ones. Then delete e and those edges intersecting it from \mathcal{H} and continue. If we can continue k steps, then we find k pairwise disjoint edges. Otherwise, we decompose \mathcal{H} into less than k families, $\mathcal{H}_1, \dots, \mathcal{H}_\ell$, of edges such that in each family \mathcal{H}_i there are at most 156 pairwise disjoint edges.

The boundary of the union of m half-spaces in \mathbb{R}^3 is the boundary of a polyhedron with at most m facets, which in turn has complexity linear in m . It now follows from Theorem 6 that each of the families \mathcal{H} has fractional Helly number 2 in a way that is independent of P , as described in the statement of Theorem 6. It is well known that families of half-spaces (in any fixed dimension) have bounded VC-dimension. Hence each \mathcal{H}_i has a bounded VC-dimension (in fact bounded by 4). Therefore, each \mathcal{H}_i has an ϵ -net of size that depends only on ϵ (see [17]). The method of Alon and Kleitman [4] implies that each \mathcal{H} satisfies a $(p, 2)$ theorem. That is, if a subset S of edges in \mathcal{H} satisfies the $(p, 2)$ property (that is, from every p sets in S there are 2 sets that intersect), then there are $c(p)$ (a constant that depends only on p) vertices that together pierce all the sets in S .

By our assumption, each \mathcal{H}_i has the $(p, 2)$ -property for $p = 157$. It follows that one can find a set of points of cardinality at most $c(157)k$ that together pierce all the edges in \mathcal{H} . ■

5 The case of half-spaces in \mathbb{R}^d where $d \geq 4$

In this section we prove Theorem 5.

For every $n \in \mathbb{N}$ we need to construct a set P of $N = \binom{n}{2}$ points and a set \mathcal{F} of n half-spaces in \mathbb{R}^4 such that:

1. Every two edges in $\mathcal{H}(P, \mathcal{F})$ have a non-empty intersection
2. Any subset of P which pierce all edges in $\mathcal{H}(P, \mathcal{F})$ must consist of at least $\frac{n-1}{2}$ points.

The next lemma will be our main tool in constructing $\mathcal{H}(P, \mathcal{F})$. This lemma is a slight variation of an argument which was used by [1] to upper bound the sign-rank of a hyper-graph.

Lemma 6. *Let \mathcal{H} be a hypergraph such that every $v \in V(\mathcal{H})$ belongs to at most d hyper-edges. Then \mathcal{H} can be realized by points and half-spaces in \mathbb{R}^{2d} . That is, \mathcal{H} is isomorphic to $\mathcal{H}(P, \mathcal{F})$ for some set P of points in \mathbb{R}^{2d} and a family \mathcal{F} of half spaces in \mathbb{R}^{2d} .*

Proof. Pick some enumeration of $E(\mathcal{H})$, e_1, e_2, \dots, e_m where $m = |E(\mathcal{H})|$. For every $v \in V$ pick some real univariate polynomial $P_v(x)$ such that

- $P_v(0) = -1$,
- $P_v(i) > 0$ if $v \in e_i$ and $P_v(i) < 0$ if $v \notin e_i$, and
- $\deg(P_v) \leq 2d$.

It is not hard to see that such a polynomial always exists: For example, the polynomial

$$P_v(x) = -\frac{Q_v(x)}{Q_v(0)}, \text{ where } Q_v(x) = \prod_{i:v \in e_i} \left(x - \left(i + \frac{1}{4}\right)\right) \left(x - \left(i - \frac{1}{4}\right)\right)$$

satisfies the above requirements. For every $v \in V$ let $p_{v,i}, i = 0, \dots, 2d$ denote the coefficients of $P_v(x)$. Notice that $p_{v,0} = -1$ for all v .

Every $v \in V$ will correspond to the point $x_v = (p_{v,1}, \dots, p_{v,2d})$ and every e_i correspond to the half-space $H_i = \{x : \langle x, n_i \rangle \geq 1\}$, where $n_i = (i, i^2, \dots, i^{2d})$. Observe that $\langle x_v, n_i \rangle = P_v(i) + 1$ and therefore $v \in e_i$ if and only if $x_v \in H_i$ as required. \square

We now construct an hyper-graph \mathcal{H} with $\binom{n}{2}$ vertices such that every vertex belongs to precisely two edges, every two edges have a non-empty intersection (that is, any matching in \mathcal{H} is of size at most 1), and finally, any set of vertices that pierces all edges must consist of at least $\frac{n-1}{2}$ vertices. Once we introduce such a hyper-graph, it follows from Lemma 6 that it can be realized in \mathbb{R}^4 by points and half-spaces.

We take the vertices of \mathcal{H} to be the edges of a complete simple graph on n vertices K_n . Let us denote the vertices of K_n by x_1, \dots, x_n . Then \mathcal{H} has $\binom{n}{2}$ vertices. The hyper-graph \mathcal{H} will consist of n edges e_1, \dots, e_n defined as follows. For every $1 \leq i \leq n$ the edge e_i is the collection of all edges in K_n incident to x_i .

It is easy to check that indeed every two sets in $\mathcal{H}(P, \mathcal{F})$ have a non-empty intersection and that any set of vertices of \mathcal{H} that pierces all the edges of \mathcal{H} must have size of at least $\frac{n-1}{2}$, as desired. \blacksquare

References

- [1] N. Alon, P. Frankl, and V. Rodl, Geometrical realization of set systems and probabilistic communication complexity. 26th Annual Symposium on Foundations of Computer Science, 1985.
- [2] N. Alon, and G. Kalai, Bounding the Piercing Number. *Discrete Comput. Geom.* (1995), volume 13, 245–256.
- [3] N. Alon, G. Kalai, J. Matoušek, and R. Meshulam, Transversal numbers for hypergraphs arising in geometry. *Advances in Applied Mathematics* (2002), volume 29, 79–101.

- [4] N. Alon, D.J. Kleitman, Piercing convex sets and the Hadwiger-Debrunner (p, q) -problem. *Adv. Math.* **96** (1992), no. 1, 103–112.
- [5] N. Alon, S. Moran, and A. Yehudayoff, Sign rank versus VC dimension. Electronic Colloquium on Computational Complexity (ECCC), 135, 2014.
- [6] P.K. Agarwal, J. Pach, M. Sharir, State of the union (of geometric objects). (English summary) Surveys on discrete and computational geometry, 9–48, Contemp. Math., 453, Amer. Math. Soc., Providence, RI, 2008.
- [7] S. Ben-David, N. Eiron, and H.-U. Simon, Limitations of learning via embeddings in Euclidean half spaces. *Journal of Machine Learning Research*, 3, pages 441–461, 2002.
- [8] H. Brönnimann and M. T. Goodrich, Almost Optimal Set Covers in Finite VC-Dimension. *Discrete Comput. Geom.* (1995), volume 14, 463–479.
- [9] S. Buzaglo, R. Pinchasi, and G. Rote, Topological hypergraphs, *Thirty Essays on Geometric Graph Theory*, J. Pach Editor, 71–81, Springer, New York, 2013.
- [10] T. M. Chan and S. Har-Peled, Approximation Algorithms for Maximum Independent Set of Pseudo-Disks. *Discrete Comput. Geom.* (2012), volume 48, 373–392. volume 95, pages 358–362, 2005.
- [11] G. Even, D. Rawitz, and S. Shahar, Hitting sets when the VC-dimension is small. In *Inf. Process. Lett.*, volume 95, pages 358–362, 2005.
- [12] J. Forster, A linear lower bound on the unbounded error probabilistic communication complexity. In *16th Annual IEEE Conference on Computational Complexity*, pages 100–106, 2001.
- [13] J. Forster, M. Krause, S. V. Lokam, R. Mubarakzjanov, N. Schmitt, and H.-U. Simon, Relations between communication complexity, linear arrangements, and computational complexity. In *Foundations of Software Technology and Theoretical Computer Science*, volume 2245 of *Lecture Notes in Computer Science*, pages 171–182, 2001.
- [14] J. Forster, N. Schmitt, H.-U. Simon, and T. Suttorp, Estimating the optimal margins of embeddings in euclidean half spaces. In *Machine Learning*, volume 51, pages 263–281, 2003.
- [15] J. Forster and H.-U. Simon, On the smallest possible dimension and the largest possible margin of linear arrangements representing given concept classes. *Theor. Comput. Sci.*, 350(1), pages 40–48, 2006.
- [16] Ch. Chojnacki (A. Hanani), Über wesentlich unplättbare Kurven im dreidimensionalen Raume, *Fund. Math.* 23 (1934), 135–142.
- [17] D. Haussler, E. Welzl, ϵ -nets and simplex range queries. *Discrete Comput. Geom.* **2** (1987), no. 2, 127–151.
- [18] K. Kedem, R. Livné, J. Pach, M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete Comput. Geom.* **1** (1986), no. 1, 59–71.
- [19] M.V. Marathe, H. Brey, H.B. Hunt III, S.S. Ravi, D.J. Rosenkrantz, Simple heuristics for unit disk graphs, *Networks* **25** (1995), 59–68.

- [20] J. Matoušek, Bounded VC-dimension implies a fractional Helly theorem. *Discrete Comput. Geom.* **31** (2004), no. 2, 251–255.
- [21] J. Matoušek, R. Seidel, E. Welzl, How to net a lot with little: Small ϵ -nets for disks and halfspaces, In Proc. 6th Annu. ACM Sympos. Comput. Geom., pages 16–22, 1990. Revised version at <http://kam.mff.cuni.cz/~matousek/enets3.ps.gz>.
- [22] J. Pach and G. Tardos, Tight lower bounds for the size of epsilon-nets. *J. Amer. Math. Soc.* **26**, pages 645–658, 2013
- [23] R. Paturi and J. Simon, Probabilistic communication complexity. *J. Comput. Syst. Sci.*, **33**(1), pages 106–123, 1986.
- [24] R. Pinchasi, A finite family of pseudo-discs must include a “small” pseudo-disc *SIAM J. Discrete Math.*, accepted. Revised version at http://www2.math.technion.ac.il/~room/ps_files/small_pseudocircle.pdf.
- [25] R. Pinchasi, A Note on Smaller Fractional Helly Numbers, *Discrete Comput. Geom.*, accepted. Revised version at http://www2.math.technion.ac.il/~room/ps_files/FracH.pdf.
- [26] A. A. Sherstov, Communication complexity under product and nonproduct distributions. *Computational Complexity*, **19**(1), pages 135–150, 2010.
- [27] 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.
- [28] W.T. Tutte, Toward a theory of crossing numbers, *J. Combinat. Theory* **8** (1970), 45–53.