

Topological Hyper-Graphs

Sarit Buzaglo* Rom Pinchasi† Günter Rote‡

December 4, 2007

Abstract

Let P be a set of n points in the plane. A *topological hyper-graph* G on the set of points of P is a collection of simple closed curves in the plane that avoid the points of P . Each of these curves is called an *edge* of G , and the points of P are called the *vertices* of G . We provide bounds on the number of edges of topological hyper-graphs in terms of the number of their vertices under various restrictions assuming the set of edges is a family of pseudo-circles.

1 Introduction

A *topological graph* is a graph drawn in the plane with its vertices drawn as points and its edges are drawn as Jordan arcs connecting corresponding points. In this paper we make an attempt to generalize the notion of a topological graph and consider topological *hyper-graphs*.

Definition 1.1. *Let C be a simple closed Jordan curve in the plane. By Jordan's Theorem C divides the plane into two regions, only one of which is bounded. We call the bounded region the disc bounded by C and we denote this region by $\text{disc}(C)$. Any point x inside $\text{disc}(C)$ is said to be surrounded by C , and C is said to be surrounding x .*

We are now ready to define a topological hyper-graph. A *topological hyper-graph* G is a pair (P, \mathcal{C}) , where P is a finite set of points in the plane, and \mathcal{C} is a family of simple closed curves in the plane that avoid the points of P . For each curve $C \in \mathcal{C}$ we denote by P_C the set of all points of P surrounded by C . C is called an *edge* of G and the set P_C is the set of *vertices* of the edge C . Two edges of G , C_1 and C_2 , are called *parallel* if $P_{C_1} = P_{C_2}$. We usually assume that any two curves from \mathcal{C} intersect in a finite number of points and that in each such intersection point the two curves properly cross. G is called *k-uniform*, if for every $C \in \mathcal{C}$ $|P_C| = k$.

In what follows and throughout the rest of the paper we will always assume that the topological hyper-graph in question **does not contain parallel edges**.

A family of simple closed curves in the plane is a family of *pseudo-circles*, if every two curves in the family are either disjoint or properly cross at precisely two points. Any collection of circles in the plane is an example for such a family.

The following simple Lemma is a crucial observation.

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

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

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

Lemma 1. *Let C_1 and C_2 be two pseudo-circles in the plane. Let x_1 and y_1 be two points that are surrounded by C_1 but not by C_2 . Let x_2 and y_2 be two points surrounded by C_2 but not by C_1 . Let e_1 be any Jordan arc connecting x_1 and y_1 that is entirely contained in the disc bounded by C_1 . Let e_2 be any Jordan arc connecting x_2 and y_2 that is entirely contained in the disc bounded by C_2 . Then e_1 and e_2 cross an even number of times.*

Proof. If C_1 and C_2 do not intersect, then $\text{disc}(C_1)$ and $\text{disc}(C_2)$ must be disjoint. Indeed, otherwise one must contain the other, say $\text{disc}(C_1)$ contains $\text{disc}(C_2)$. This is a contradiction to the assumption that x_2 and y_2 are two points surrounded by C_2 but not by C_1 . However, if $\text{disc}(C_1)$ is disjoint from $\text{disc}(C_2)$, then clearly e_1 and e_2 cannot intersect.




Figure 1: Lemma 1

Therefore, assume that C_1 and C_2 intersect in two points u and v . Let S_1 be the subarc of C_1 between u and v that is contained in $\text{disc}(C_2)$. Let S_2 be the subarc of C_2 between u and v that is contained in $\text{disc}(C_1)$. S_1 together with S_2 combine a simple closed curve that surrounds exactly those points in $\text{disc}(C_1) \cap \text{disc}(C_2)$. Let Z denote the region $\text{disc}(C_1) \cap \text{disc}(C_2)$. e_1 and e_2 may intersect only at a point of Z (see Figure 1). Since both x_1 and y_1 are not in Z , as they are not surrounded by C_2 , the intersection of e_1 with Z consists of a finite number of arcs each of which has both endpoints lying on S_2 . We will show that each of these arcs crosses e_2 an even number of times. Let g be such an arc whose endpoints on S_2 are k and l . Let S'_2 be the subarc of S_2 between k and l . g and S'_2 combine a closed curve that is contained in Z . Therefore both x_2 and y_2 are not surrounded by this closed curve. It follows that e_2 the Jordan arc connecting x_2 and y_2 intersects $g \cup S'_2$ an even number of times. However, e_2 does not intersect S'_2 as it is a subarc of C_2 . Therefore, e_2 crosses g an even number of times.

It follows now that e_2 intersects e_1 an even number of times which is what we wanted to prove. ■

The next theorem is a simple consequence of Lemma 1.

Theorem 2. *Let $G = (P, \mathcal{C})$ be a topological 2-uniform hyper-graph on n vertices. If \mathcal{C} is*

a family of pseudo-circles, then G , viewed as an abstract 2-graph, is a planar graph. In particular, there are at most $3n - 6$ curves in \mathcal{C} .

Proof. We draw the graph G as a topological graph in such a way that every pair of edges that are not incident to the same vertex cross an even number of times. Then it is a consequence of a theorem of Hanani and Tutte ([2, 9]) that G is planar. The drawing rule for G is as follows: For any curve $C \in \mathcal{C}$ let x and y be the two points of P surrounded by C . We draw a Jordan arc connecting x to y so that the arc is contained in $\text{disc}(C)$.

It is left to show that the above drawing rule provides the desired topological graph. Let e and f be two edges in our drawing that do not share a common vertex. Let x and y be the vertices of e and let $C_e \in \mathcal{C}$ be the curve in \mathcal{C} that surrounds both x and y . Let a and b be the vertices of f and let $C_f \in \mathcal{C}$ be the curve in \mathcal{C} that surrounds both a and b . Since C_e surrounds x and y and no other point, then in particular it does not surround a or b . Similarly, C_f surrounds a and b but none of x and y . Therefore, by Lemma 1, e and f cross an even number of times. ■

We will now try to bound the number of edges in a topological hyper-graph in more general settings. The following lemma can be found in [1] (see also [6]).

Lemma 3. *Let A_1, \dots, A_n be connected sets in the plane each of which is also simply connected. If for every $1 \leq i < j \leq n$ $A_i \cap A_j$ is connected, then $A_1 \cap \dots \cap A_n$ is either connected or empty.*

As a consequence of Lemma 3 we obtain the following corollary:

Corollary 4. *If C_1, \dots, C_t is any collection of pseudo-circles, then $\text{disc}(C_1) \cap \dots \cap \text{disc}(C_t)$ is a connected set, unless it is empty.*

Indeed, this is because $\text{disc}(C_i)$ is a simply connected set for every C_i , and $\text{disc}(C_i) \cap \text{disc}(C_j)$ is either empty or connected for every C_i and C_j .

The next theorem provides a linear bound on the number of edges of a 3-uniform topological hyper-graph $G = (P, \mathcal{C})$, provided that \mathcal{C} is a family of pseudo-circles.

Theorem 5. *Let $G = (P, \mathcal{C})$ be a 3-uniform topological graph on n vertices such that \mathcal{C} is a family of pseudo-circles. Then $|\mathcal{C}| = O(n)$.*

Proof. We draw a topological graph H whose vertices are the points of P and whose edges are all pairs of points that are surrounded by some curve in \mathcal{C} . Let x and y be two points of P surrounded by a curve in \mathcal{C} , and let C_1, \dots, C_l be all the curves in \mathcal{C} that surround both x and y . By Corollary 4, $Z = \text{disc}(C_1) \cap \dots \cap \text{disc}(C_l)$ is a connected set. We draw the edge in H between x and y as a Jordan curve inside the connected region Z .

For every edge e in H , let $d(e)$ denote the number of curves in \mathcal{C} that surround both endpoints of e (and therefore also the entire edge e).

Claim 6. *Let e and f be two edges of H with no common vertex. If e and f cross an odd number of times, then either $d(e) \leq 2$, or $d(f) \leq 2$, or both.*

Proof. Assume to the contrary that both $d(e)$ and $d(f)$ are greater than or equal to 3. Since $d(e) \geq 3$, there are at least three different curves from \mathcal{C} that surround both endpoints of e . Therefore, there is a curve $C_e \in \mathcal{C}$ that surrounds both endpoints of e but no endpoint of f . Similarly, there is a curve $C_f \in \mathcal{C}$ that surrounds both endpoints of f but no endpoint of e . By Lemma 1, e and f cross an even number of times, a contradiction. ■

Let E' be the set of all edges e of H such that $d(e) \leq 2$. By Claim 6, the set of edges in $E \setminus E'$ forms a planar graph, as every two edges in $E \setminus E'$ that do not share a common vertex cross an even number of times. In particular the cardinality of $E \setminus E'$ is linear in n . We will now show that also the cardinality of E' is linear in n .

In fact we will prove even a slightly stronger statement that will be helpful for the proof of Theorem 5.

Claim 7. *Let V be any subset of the vertices of H and let $E(V)$ be those edges in E' both of whose vertices are in V . Then $|E(V)| \leq c|V|$, where $c > 0$ is some absolute constant.*

Proof. Let $0 < q < 1$ be a positive number to be determined later. Pick every vertex of V with probability q and consider only those edges of $E(V)$ both of whose vertices were picked. We thus obtain a random subgraph of H , that we denote by \tilde{H} , on a set of vertices \tilde{V} . Let e be an edge in \tilde{H} and let x and y be its two vertices. We know that $d(e) \leq 2$. Therefore, there exist at most two vertices $z \in V$ such that x, y , and z are surrounded by some curve $C \in \mathcal{C}$. We say that e is *good*, if there is no such vertex z in \tilde{V} .

Observe that every two good edges with no common vertex cross an even number of times. Indeed, if e and f are two good edges with no common vertex, then there is a curve $C_e \in \mathcal{C}$ that surrounds both vertices of e but no vertex of f (or else e would not be good). And similarly, there is a curve $C_f \in \mathcal{C}$ that surrounds both vertices of f but no vertex of e . Hence, by Lemma 1, e and f cross an even number of times. It follows now from the theorem of Hanani and Tutte [2, 9] that the good edges in \tilde{H} constitute a planar graph. Therefore, the number of good edges in \tilde{H} is at most $3|\tilde{V}|$. This inequality is true also for the expected values of the number of good edges and $|\tilde{V}|$. Clearly, $Ex(|\tilde{V}|) = q|V|$. As for the expected number of good edges, observe that any edge $e \in E(V)$ is a good edge in \tilde{H} with probability of at least $q^2(1-q)^2$. Indeed, this is because of the possible event in which we pick the two vertices of e to \tilde{V} and we do not pick those vertices in V (there are at most two such vertices) that are surrounded together with the two endpoints of e by some curve in \mathcal{C} .

It follows now that $q^2(1-q)^2|E(V)| \leq 3q|V|$. Taking $q = \frac{1}{3}$ we obtain $|E(V)| \leq 21|V|$. ■

Considering the graph H again, we deduce that there is an absolute constant c' such that the number of edges of H induced by any subset V of vertices of H is at most $c'|V|$. Indeed, this is true for the subset E' of edges of H , by Claim 7, and also to the subset of edges $E \setminus E'$, as they constitute a planar graph.

We now show that the number of triangles (that is, cycles of length 3) in the graph H is at most linear in the number of vertices of H . This will be enough to prove Theorem 5 because every pseudo-circle $C \in \mathcal{C}$ corresponds to precisely one triangle in the graph G (the triangle whose vertices are surrounded by C). We prove a more general lemma on abstract graphs:

Lemma 8. *Let H be a graph on m vertices. Assume that there is an absolute constant c' such that the number of edges of H induced by any subset V of vertices of H is at most $c'|V|$. Then for every $l \geq 2$ the number of copies of K_l (the complete graph on l vertices) in H is at most $c_l m$, where $c_l = \frac{(2c')^{l-1}}{l}$.*

Proof. We prove the lemma by induction on l . For $l = 2$ there is nothing to prove, as this is the assumption in the Lemma. Assume the lemma is true for l . Let v_1, \dots, v_m denote the vertices of H and let E denote the set of edges of H . We know that $|E| \leq c'm$. For every $1 \leq i \leq m$, let d_i denote the degree of v_i in H . Fix i and consider the number of copies of K_{l+1} in H that include the vertex v_i . Each such copy of K_{l+1} correspond to a unique copy of K_l among the neighbors of v_i . Therefore, it follows from the induction hypothesis, applied

on the subgraph H_{v_i} of H induced by the neighbors of v_i , that the number of copies of K_l in H_{v_i} is at most $c_l d_i$. It follows now that that v_i is incident to at most $c_l d_i$ copies of K_{l+1} in H . Summing over all i between 1 and m , we see that the number of times a vertex in H is incident to a copy of K_{l+1} in H is at most $\sum_{i=1}^m c_l d_i = 2c_l |E| \leq 2c_l c' m$. On the other hand, if Z denotes the number of copies of K_{l+1} in H , then the number of times a vertex in H is incident to a copy of K_{l+1} in H is exactly $(l+1)Z$. It follows that $Z \leq \frac{2c_l c'}{l+1} m = c_{l+1} m$ ■

In particular the number of triangles in our topological graph H is at most cn where c is some absolute constant independent of n . This concludes the proof of Theorem 5. ■

If $G = (P, \mathcal{C})$ is a topological hyper-graph and \mathcal{C} is a family of pseudo-circles but we do not assume that G is uniform, then a linear bound on the number of edges of G is no longer valid. Nevertheless, it is possible to obtain a cubic (that is also tight) bound on the number of edges in G in this case, as we shall now see.

Recall that for a family \mathcal{F} of sets the *VC-dimension* ([10]) of \mathcal{F} is the largest cardinality of a set S that \mathcal{F} *shatters*. That is, the largest cardinality of a set S such that for any subset B of S there exists $F \in \mathcal{F}$ with $B = F \cap S$.

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

Theorem 9. *Let $G = (P, \mathcal{C})$ be a topological hyper-graph on n vertices such that \mathcal{C} is a family of pseudo-circles. Then the family $\mathcal{F} = \{P_C \mid C \in \mathcal{C}\}$ has VC-dimension at most 3. In particular, by the Perles-Sauer-Shelah theorem, $|\mathcal{C}| = O(n^3)$.*

Proof. We show that \mathcal{F} cannot shatter any set of 4 points. Assume to the contrary that it does, and let $\{p_1, p_2, p_3, p_4\} \subset P$ be a set of four points shattered by \mathcal{F} . For every $1 \leq i < j \leq 4$ consider the family \mathcal{C}_{ij} of all the curves in \mathcal{C} that surround both p_i and p_j . By Corollary 4, the set $R_{ij} = \bigcap_{C \in \mathcal{C}_{ij}} \text{disc}(C)$ is a connected set. We draw an edge (Jordan arc) e_{ij} between p_i and p_j inside the region R_{ij} . We thus get a drawing of K_4 in the plane. We claim that in this drawing every two edges that do not share a common vertex cross an even number of times. Indeed, consider the edge e_{ij} between P_i and P_j and the edge e_{kl} between p_k and p_l (we assume that $\{i, j, k, l\} = \{1, 2, 3, 4\}$). As $\{p_1, p_2, p_3, p_4\}$ is shattered by \mathcal{F} , there is a curve $C_{ij} \in \mathcal{C}_{ij}$ that surrounds both p_i and p_j but neither p_k nor p_l . Similarly, there is a curve $C_{kl} \in \mathcal{C}_{kl}$ that surrounds both p_k and p_l but neither p_i nor p_j . By our construction, $e_{ij} \subset \text{disc}(C_{ij})$ and $e_{kl} \subset \text{disc}(C_{kl})$. It follows now from Lemma 1 that e_{ij} and e_{kl} cross an even number of times.

It is easy to observe that by small modifications of the drawing of the edges going from each p_i only in a small neighborhood of p_i , we may obtain a new drawing in which also every two edges that share a common vertex cross an even number of times (apart from meeting at the same vertex).

For a closed curve S in the plane, not necessarily simple, and a point x not on S we say that x is *surrounded* by S , if the index of the curve S with respect to the point x is odd. This is equivalent to that any ray (or just a Jordan curve that goes to infinity) emanating from x meets S (as a curve, not as a set) an odd number of times. Observe that this definition generalizes the notion of a point surrounded by a simple closed curve.

Lemma 10. *In any drawing of K_4 in the plane in which every two edges cross an even number of times, there is a vertex v of K_4 that is surrounded by the closed curve composed from the three edges of K_4 not incident to v .*

Proof. Assume not and once again let p_1, p_2, p_3 , and p_4 denote the vertices of K_4 drawn in the plane. For any three vertices p_i, p_j , and p_k denote by S_{ijk} the closed curve composed from the edges $p_i p_j$, $p_j p_k$, and $p_k p_i$. Consider a fixed vertex p_i . Any small enough circular disc D , centered at p_i , is trisected by the edges going from p_i to the three other vertices. Fix such a disc D . We claim the section of D bounded by the edges $p_i p_j$ and $p_i p_k$ consists of points that are all surrounded by S_{ijk} . Indeed, otherwise the portion of the edge $p_i p_l$ inside D is surrounded by S_{ijk} . As the edge $p_i p_l$ crosses each of the edges $p_i p_j$, $p_i p_k$, and $p_j p_k$ an even number of times, it follows that p_l is also surrounded by S_{ijk} , a contradiction.

Consider now a point x far away, not surrounded by any of the curves S_{ijk} . Draw an arc g connecting x to p_1 . Once again let D be a small circular disc centered at p_1 . D is trisected by the edges $p_1 p_i$ for $i = 2, 3, 4$. Without loss of generality assume that $g \cap D$ lies in the portion of D delimited by $p_1 p_2$ and $p_1 p_3$. Hence, as x is not surrounded by S_{123} , it follows that g crosses S_{123} an odd number of times. For the same reasons it follows that g crosses both S_{124} and S_{134} an *even* number of times because $g \cap D$ is not surrounded by S_{124} nor by S_{134} .

Denote by t_{ij} the number of crossings between g and the edge $p_i p_j$. Therefore, $t_{12} + t_{13} + t_{23}$ is odd and both $t_{12} + t_{14} + t_{24}$ and $t_{13} + t_{14} + t_{34}$ are even. Summing them all up we conclude that $t_{23} + t_{34} + t_{24}$ is odd. Therefore, g crosses S_{234} an odd number of times. Now since x is not surrounded by S_{234} and g connects x to p_1 , it follows that p_1 is surrounded by S_{234} , a contradiction. ■

In view of Lemma 10, assume without loss of generality that p_4 is surrounded by the closed curve composed from the edges $p_1 p_2$, $p_2 p_3$, and $p_3 p_1$. As \mathcal{F} shatters $\{p_1, p_2, p_3, p_4\}$, let $C \in \mathcal{C}$ be a curve surrounding p_1, p_2 , and p_3 but not p_4 . Therefore, each of the edge $p_1 p_2$, $p_2 p_3$, and $p_3 p_1$ is contained in $\text{disc}(C)$. This is a contradiction because now it is not possible that p_4 is surrounded by the closed curve composed from the edges $p_1 p_2$, $p_2 p_3$, and $p_3 p_1$, as it is not surrounded by C . ■

For a set A and an integer $r \geq 0$ we denote by $\binom{[A]}{\leq r}$ the family of all subsets of A of cardinality at most r . We will need the following small variant of the original Perles-Sauer-Shelah theorem ([7]).

Theorem 11. *Let $\mathcal{F} = \{A_1, \dots, A_m\}$ be a family of subsets of $\{1, 2, \dots, n\}$ and assume that \mathcal{F} has VC-dimension less than or equal to d , then*

$$m \leq \left| \bigcup_{i=1}^m \binom{[A_i]}{\leq d} \right|.$$

Proof. We prove the theorem by induction on d and n . The theorem is clearly true for $d = 0$ and any n . For $d > 0$, assume that \mathcal{F} has VC-dimension at most d and define $\mathcal{F}_1 = \{A \setminus \{1\} \mid A \in \mathcal{F}\}$. Let $\mathcal{T} = \{A \setminus \{1\} \mid 1 \in A \text{ and } A, A \setminus \{1\} \in \mathcal{F}\}$. It is easy to see that $|\mathcal{F}| = |\mathcal{F}_1| + |\mathcal{T}|$. Let s denote the number of sets in \mathcal{F}_1 and let t denote the number of sets in \mathcal{T} . We rewrite $\mathcal{F}_1 = \{F_1, \dots, F_s\}$ and $\mathcal{T} = \{T_1, \dots, T_t\}$. As the VC-dimension of \mathcal{F}_1 is at most d , we use the induction hypothesis on n to deduce that the family $\bigcup_{i=1}^s \binom{[F_i]}{\leq d}$ contains at least s sets.

Observe that the VC-dimension of \mathcal{T} is at most $d-1$, for otherwise \mathcal{F} has VC-dimension at least $d+1$. Therefore, by the induction hypothesis, there are at least t sets in $\bigcup_{i=1}^t \binom{[T_i]}{\leq d-1}$. To each of these sets add the element 1 to obtain t sets in $\bigcup_{i=1}^m \binom{[A_i]}{\leq d}$ each containing the element 1 and thus each is different than any set in $\bigcup_{i=1}^s \binom{[F_i]}{\leq d}$. Keeping in mind that $m = s + t$ we have proved the desired result. ■

Corollary 12. *Let $\mathcal{F} = \{A_1, \dots, A_m\}$ be a family of subsets of $\{1, 2, \dots, n\}$ with VC-dimension less than or equal to d . Then it is possible to assign to each set A_i a subset of its of size at most d such that no two sets A_i and A_j are assigned with the same set of at most d elements.*

Proof. Consider the bipartite graph in which one set of vertices corresponds to the sets A_1, \dots, A_m and the other set of vertices corresponds to the elements in $\bigcup_{i=1}^m \binom{[A_i]}{\leq d}$. In this bipartite graph connect each set A_i to all its subsets of size at most d . In order to find a matching that includes all the sets A_1, \dots, A_m apply Hall's theorem. By Theorem 11, the condition in Hall's theorem is satisfied. ■

Theorems 2 and 5 give a linear bound to the number of edges of a 2-uniform and 3-uniform, respectively, topological hyper-graphs $G = (P, \mathcal{C})$ where \mathcal{C} is a family of pseudo-circles. The next theorem provides a bound to the number of edges in a topological hyper-graph $G = (P, \mathcal{C})$ on n vertices assuming that \mathcal{C} is a family of pseudo-circles and that $|P_C| \leq k$ for every $C \in \mathcal{C}$ and a fixed $k > 3$.

Theorem 13. *Let $G = (P, \mathcal{C})$ be a topological hyper-graph on n vertices. Assume that \mathcal{C} is a family of pseudo-circles and that $|P_C| \leq k$ for every $C \in \mathcal{C}$ and a fixed integer k . Then $|\mathcal{C}| = O(k^2 n)$.*

Proof. Consider the family $\mathcal{F} = \{P_C \mid C \in \mathcal{C}\}$. By Theorem 9, \mathcal{F} has VC-dimension that is less than or equal to 3. We use Corollary 12 and assign to each member P_C of \mathcal{F} a subset of its of size at most 3 that we denote by B_C .

We fix a small number $0 < q < 1/2$ to be determined later and pick every point of P independently with probability q . We call a curve $C \in \mathcal{C}$ *good*, if every point in B_C was picked and no point in $P_C \setminus B_C$ was picked. Observe that a curve C in \mathcal{C} is good with probability $q^{|B_C|}(1-q)^{|P_C|-|B_C|} \geq q^{|B_C|}(1-q)^{k-|B_C|} \geq q^3(1-q)^{k-3}$, where the last inequality is because $q < 1/2$.

By Theorem 2 and Theorem 5 the number of edges in a 2-uniform, or a 3-uniform topological hyper-graph on n vertices whose set of edges is a family of pseudo-circles is $O(n)$. Therefore, the number of good curves in \mathcal{C} is at most some absolute constant c times the number of points of P that were picked. Taking expectation we see that $|\mathcal{C}|q^3(1-q)^{k-3} \leq cqn$. Taking $q = 1/k$ we get the desired result, namely, $|\mathcal{C}| \leq c'k^2n$ for some absolute constant $c' > 0$. ■

Remark. It is not hard to show that the bound in Theorem 13 can indeed be attained even by a family \mathcal{C} of circles in the plane.

Definition 1.2. *We denote by $f_t(n)$ the maximum number of edges in a topological graph on n vertices with no t vertex-disjoint edges every two of which cross an odd number of times.*

We will need the following rough estimate on $f_t(n)$:

Lemma 14. $f_t(n) = O(n^{2-1/2^{t-2}})$.

Proof. Let K be a topological graph on n vertices with no t vertex-disjoint edges every two of which cross an odd number of times. If $t = 2$, then K is a planar graph by the theorem of Hanani and Tutte. It follows that $f_2(n) = 3n - 6$ edges.

Assume the lemma is true for t and consider $f_{t+1}(n)$. Let K be a topological graph on n vertices with no $t+1$ edges every two of which cross an odd number of times. For each edge e

of K consider the set $Odd(e)$ of edges that cross the edge e an odd number of times. Clearly, in $Odd(e)$ there are no t edges every two of which cross an odd number of times, for otherwise together with e they would constitute $t + 1$ such edges. Therefore, $Odd(e) \leq f_t(n)$. Hence, the number of pairs of edges of G that cross an odd number of times is at most $\frac{1}{2}|E|f_t(n)$. On the other hand by the Crossing Lemma ([4, 5]), we know that this number is at least $\frac{c|E|^3}{n^2}$ for some constant $c > 0$, assuming that $|E|$ is at least $4n$. Therefore, $|E| \leq \frac{1}{\sqrt{2c}}n\sqrt{f_t(n)}$. In particular, $f_{t+1}(n) \leq \frac{1}{\sqrt{2c}}n\sqrt{f_t(n)}$. Solving this recurrence we obtain: $f_t \leq c_t n^{2^{-1/2^{t-2}}}$, where c_t is some absolute constant that depends only on t . ■

The next theorem provides even better bounds on the size of $\{P_C \mid C \in \mathcal{C}\}$, assuming that the intersections of the sets P_C with each other are 'small'.

Theorem 15. *Let $r, k_1, k_2 > 0$ be fixed integers. Let $G = (P, \mathcal{C})$ be a topological hyper-graph on n vertices and assume that \mathcal{C} is a family of pseudo-circles. Assume further that $k_1 \leq |P_C| \leq k_2$ for every $C \in \mathcal{C}$. If for every r curves $C_1, \dots, C_r \in \mathcal{C}$ we have $|P_{C_1} \cap \dots \cap P_{C_r}| < r$, then $\sum_{C \in \mathcal{C}} |P_C| = O(\frac{k_2}{k_1^2}n)$, where the constant of proportionality depends only on r .*

Proof. We define a topological graph H in the following way: The vertices of H are the points of P . For any pair of vertices x and y that are surrounded by some curve $C \in \mathcal{C}$, consider all curves in \mathcal{C} that surround x and y . Let those curves be C_1, \dots, C_l and define $Z = \text{disc}(C_1) \cap \dots \cap \text{disc}(C_l)$. By Corollary 4, Z is a connected set. We draw the edge in H between x and y as a Jordan curve entirely included in Z .

We call an edge of H *bad* if its endpoints are surrounded by at least $(r-1)2^{r-2} + 1$ different curves in \mathcal{C} , otherwise it is called *good*.

Claim 16. *There are no r^2 bad edges in H every two of which cross an odd number of times and no two of which share a common vertex.*

Proof. Assume to the contrary that E is a collection of r^2 bad edges every two of which cross an odd number of times and no two of which share a common vertex.

Let e, e_1, \dots, e_{r-2} be any $r-1$ different edges in E . We claim that among all curves in \mathcal{C} surrounding e there is at least one curve that does not surround any of the two vertices of some edge among e_1, \dots, e_{r-2} . Indeed, otherwise any curve in \mathcal{C} that surrounds e must surround at least one vertex of each of e_1, \dots, e_{r-2} . Since e is a bad edge there are at least $(r-1)2^{r-2} + 1$ curves in \mathcal{C} that surround e . Therefore, by the pigeon-hole principle, there must be at least r curves surrounding e as well as the same set of $r-2$ vertices (each of which is the endpoint of some edge among e_1, \dots, e_{r-2}). This is a contradiction to the assumption that no r curves in \mathcal{C} surround the same set of r vertices.

It follows that apart from at most $r-3$ edges in E , for every edge f in E , different from e , there is a curve $C \in \mathcal{C}$ which surrounds e but not any of the two vertices of f . This way, assuming that $E > (r-1)(r-2)$, we can find $r-1$ edges e_1, \dots, e_{r-1} with the property that for every $1 \leq i < j \leq r-1$ there is a curve in \mathcal{C} which surrounds e_i but not any of the two vertices of e_j .

Using once again the argument above, there is a curve $C \in \mathcal{C}$ and there is an index j between 1 and $r-2$ such that C surrounds e_{r-1} but C does not surround any of the vertices of e_j . Since $j < r-1$ there is a curve $C' \in \mathcal{C}$ that surrounds e_j but that does not surround any of the vertices of e_{r-1} . By Lemma 1, this implies that e_j and e_{r-1} cross an even number of times, contradicting our assumption. ■

It follows from Claim 16 and Lemma 14 that there exists k_0 which depends only on r such that if $k_1 > k_0$ then for every $C \in \mathcal{C}$, C surrounds at least $k_1^2/4$ good edges.

Clearly, we may assume that $k_1 > k_0$ as by Theorem 13, the collection \mathcal{C}' of all curves $C \in \mathcal{C}$ with $|P_C| \leq k_0$ consists of $O(k_0^2 n)$ curves and therefore, $\sum_{C \in \mathcal{C}'} |C| = O(n)$.

Fix a probability $0 < q < 1$, to be determined later, and pick every point in P with probability q . Let H^* be the random topological graph whose vertices are those points that were picked, and whose edges are those good edges of H both of whose vertices were picked. Call an edge e of H^* *nice* if there is a curve $C \in \mathcal{C}$ that surrounds no other vertex of H^* but the endpoints of e .

We claim that the subgraph of H^* which consists only of the nice edges, is planar. Indeed, let e and f be two nice edges that do not share a common vertex. Because e is nice, there is a curve $C_e \in \mathcal{C}$ which surrounds e but does not surround any of the vertices of f . Similarly, there is a curve $C_f \in \mathcal{C}$ that surrounds f but does not surround any of the vertices of e . By Lemma 1, e and f cross an even number of times. Therefore, by the Hanani-Tutte Theorem, H^* is planar. It follows that $|E^*|$ is less than or equal to $3|V^*| - 6$, where E^* is the set of all nice edges and V^* is the set of vertices in H^* . This is true also when considering the expected values of $|E^*|$ and $|V^*|$.

We have $Ex(|V^*|) = qn$. We claim that $Ex(|E^*|) \geq |E|q^2(1-q)^{k_2-2}$, where E is the set of all good edges. Indeed, this follows because each edge $e \in E$ may become nice if its two vertices are picked to V^* (and that happens with probability q^2), and the points surrounded by C_e (an arbitrary curve in \mathcal{C} surrounding e), that are not vertices of e are not picked (this happens with probability of at least $(1-q)^{k_2-2}$ and independently of the event in which the two vertices of e are picked). Therefore, any edge in E is nice with probability of at least $q^2(1-q)^{k_2-2}$.

Therefore, $|E|q^2(1-q)^{k_2-2} < 3qn$. Taking $q = \frac{1}{k_2-1}$ we obtain, $|E| < 9(k_2-1)n$. On the other hand each good edge is surrounded by at most $(r-1)2^{r-1}$ curves in \mathcal{C} and each curve in \mathcal{C} surrounds at least $k_1^2/4$ good edges. Therefore, $|E| \geq \frac{k_1^2}{4(r-1)2^{r-1}}|\mathcal{C}|$. Hence, $|\mathcal{C}| < \frac{36(r-1)2^{r-1}k_2}{k_1^2}n$. The theorem now follows as $\sum_{C \in \mathcal{C}} |P_C| \leq \sum_{C \in \mathcal{C}} k_2 \leq 36(r-1)2^{r-1}\frac{k_2}{k_1^2}n$. ■

Corollary 17. *Let $r > 0$ be a fixed integer. Let $G = (P, \mathcal{C})$ be a topological hyper-graph on n vertices and assume that \mathcal{C} is a family of pseudo-circles. If for every r curves $C_1, \dots, C_r \in \mathcal{C}$ we have $|P_{C_1} \cap \dots \cap P_{C_r}| < r$, then $\sum_{C \in \mathcal{C}} |P_C| = O(n \log n)$.*

Proof. For every $0 \leq i \leq \log_2 n$ let \mathcal{C}_i denote those curves $C \in \mathcal{C}$ for which $2^i \leq |P_C| \leq 2^{i+1}$. By Theorem 15, for every $0 \leq i \leq \log_2 n$ we have $\sum_{C \in \mathcal{C}_i} |P_C| = O((\frac{2^{i+1}}{2^i})^2 n) = O(n)$. Therefore, $\sum_{C \in \mathcal{C}} |P_C| = \sum_{i=0}^{\lfloor \log_2 n \rfloor} O(n) = O(n \log n)$. ■

As an example, fix an integer $r \geq 2$. Let $G = (P, \mathcal{C})$ be a k -uniform topological hyper-graph on n vertices. Assume that for any two curves $C_1, C_2 \in \mathcal{C}$ we have $|P_{C_1} \cap P_{C_2}| < r$. Then clearly, for every r curves $C_1, \dots, C_r \in \mathcal{C}$, we have $|P_{C_1} \cap \dots \cap P_{C_r}| < r$. It follows now from Theorem 15 that

$$k|\mathcal{C}| = \sum_{C \in \mathcal{C}} |P_C| = O(\frac{k^2}{k^2}n) = O(n).$$

Hence, $|\mathcal{C}| = O(n/k)$. This roughly says that the sets in the family $\{P_C \mid C \in \mathcal{C}\}$, each having cardinality k , behave almost as if they were disjoint.

Acknowledgments: We thank Eyal Ackerman, Kevin Buchin, and Christian Kanuer for very helpful discussions.

References

- [1] P. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir, and S. Smorodinsky, Lenses in Arrangements of Pseudocircles and their Applications, *J. ACM*, **51**, (2004), 139–186.
- [2] Ch. Chojnacki (A. Hanani), Über wesentlich unplättbare Kurven im dreidimensionalen Raume, *Fund. Math.* **23** (1934), 135–142.
- [3] E. Helly, Über Systeme abgeschlossener Mengen mit gemeinschaftlichen Punkten, *Monatshefte d. Mathematik* **37** (1930), 281–302.
- [4] F. T. Leighton, *Complexity issues in VLSI, Foundations of Computing Series*, MIT Press, Cambridge, Mass., 1983.
- [5] F. T. Leighton, New lower bound techniques for VLSI, *Math. Systems Theory* **17** (1984), 47–70.
- [6] J. Molnár, Über eine Verallgemeinerung auf die Kugelfläche eines topologischen Satzes von Helly, *Acta Math. Acad. Sci.* **7** (1956), 107–108.
- [7] N. Sauer, On the density of families of sets, *Journal of Combinatorial Theory, Series A* **25** (1972), 80–83.
- [8] S. Shelah, A combinatorial problem, stability and order for models and theories in infinite languages, *Pacific J. Math.* **41** (1972), 247–261.
- [9] W.T. Tutte, Toward a theory of crossing numbers, *J. Combinat. Theory* **8** (1970), 45–53.
- [10] V.N. Vapnik and A. Ya. Chervonenkis, On the uniform convergence of relative frequencies of events to their probabilities, *Theory Probab. Appl.* **16** (1971), 264–280.