

On Planar Intersection Graphs with Forbidden Subgraphs*

János Pach[†]

Micha Sharir[‡]

April 6, 2006

Abstract

Let \mathcal{C} be a family of n compact connected sets in the plane, whose intersection graph $G(\mathcal{C})$ has no complete bipartite subgraph with k vertices in each of its classes. Then $G(\mathcal{C})$ has at most n times a polylogarithmic number of edges, where the exponent of the logarithmic factor depends on k . In the case where \mathcal{C} consists of convex sets, we improve this bound to $O(n \log n)$. If in addition $k = 2$, the bound can be further improved to $O(n)$.

1 Introduction

Given a collection $\mathcal{C} = \{C_1, \dots, C_n\}$ of compact connected sets in the plane, let $G = G(\mathcal{C})$ denote its *intersection graph*. That is, let $V(G) = \{C_1, \dots, C_n\}$, and connect two vertices C_i and C_j ($i \neq j$) by an edge of G if and only if $C_i \cap C_j \neq \emptyset$. In this note we investigate how many edges $G(\mathcal{C})$ can have if it has no subgraph isomorphic to a fixed graph H .

If H is not bipartite, then the assumption that G is an intersection graph does not drastically change the answer to this question. According to the Erdős–Stone theorem [4], the maximum number of edges that *any* H -free graph of n vertices can have is

$$\text{ex}(n, H) = \left(1 - \frac{1}{\chi(H) - 1} + o(1)\right) \frac{n^2}{2}.$$

Here $\chi(H)$ stands for the chromatic number of H . If H is not bipartite ($\chi(H) > 2$), then this bound is asymptotically tight, as is shown by a complete $(\chi(H) - 1)$ -partite graph whose vertex classes are of roughly the same size. Clearly, this graph can be obtained as the intersection graph of plane convex bodies, by representing the elements of each vertex class by very long pairwise disjoint rectangles with parallel sides.

*This work was supported by a joint NSF grant CCF-05-14079 and by a joint grant from the U.S.-Israel Binational Science Foundation. Work by M. Sharir was also supported by grant 155/05 from the Israel Science Fund, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. Work by J. Pach was also supported by grants from the Hungarian Science Foundation OTKA and by PSC-CUNY.

[†]Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. pach@cims.nyu.edu

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. michas@post.tau.ac.il

The problem becomes more interesting for bipartite graphs. By the Kővári–Sós–Turán theorem [5], the maximum number of edges of a graph with n vertices containing no complete bipartite graph $K_{k,k}$ with a constant number k of vertices in each of its classes, satisfies

$$\text{ex}(n, K_{k,k}) \leq b_k n^{2-\frac{1}{k}},$$

where b_k is a suitable constant. Restricting our attention to intersection graphs of compact connected sets, this bound can be substantially strengthened.

Theorem 1. *If the intersection graph of n compact connected sets in the plane has no subgraph isomorphic to $K_{k,k}$, then its number of edges cannot exceed $c(k)n \log^{e(k)} n$, for appropriate constants $c(k)$, $e(k)$ that depend on k .*

In the lack of examples of collections \mathcal{C} whose intersection graphs are $K_{k,k}$ -free and have a superlinear number of edges (in $|\mathcal{C}|$), we suspect that Theorem 1 can be substantially improved. In fact, in the special case when \mathcal{C} consists of convex sets, we can replace the polylogarithmic factor by $\log n$.

Theorem 2. *If the intersection graph of n convex sets in the plane has no subgraph isomorphic to $K_{k,k}$, then its number of edges is at most $O(n \log n)$, where the constant of proportionality depends on k .*

For $k = 2$, that is, when $G(\mathcal{C})$ contains no C_4 (cycle of length four) as a subgraph, we can completely get rid of the logarithmic factor in Theorem 2.

Theorem 3. *If the intersection graph of n convex sets in the plane has no subgraph isomorphic to C_4 , then its number of edges is at most $O(n)$.*

Since every graph can be obtained as the intersection graph of *three-dimensional* convex bodies [3], we cannot expect that a similar phenomenon holds in higher dimensions.

A graph drawn in the plane with possibly crossing straight-line edges is called a *geometric graph*. We assume for simplicity that no three vertices of a geometric graph are on a line. Two sets of edges $\{e_1, \dots, e_k\}$ and $\{f_1, \dots, f_k\}$ in a geometric graph are said to form a $k \times k$ *grid* if every e_i crosses all f_j .

Theorem 3 has the following corollary, established first in [7] with a different proof.

Corollary 4. [7] *Any geometric graph with n vertices that contains no 2×2 grid has at most $O(n)$ edges.*

2 Intersection graphs of connected sets: Proof of Theorem 1

Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be a collection of compact connected sets in the plane, and assume without loss of generality that each C_i is equal to the closure of its interior $\text{int } C_i$. Otherwise, we can replace C_i by a slightly larger set satisfying this condition, without changing the intersection pattern of \mathcal{C} .

Fix distinct points $p_i \in \text{int } C_i$, for $i = 1, \dots, n$. For any $i < j$ with $C_i \cap C_j \neq \emptyset$, connect p_i to p_j by a simple (non-selfintersecting) continuous arc γ_{ij} that does not pass through any other point p_h ($h \neq i, j$). We can easily achieve that




Figure 1: The drawing of G^* .

1. γ_{ij} is the union of two simple arcs γ_{ij}^0 and γ_{ij}^1 that share only one (end)point such that $\gamma_{ij}^0 \subset C_i$ and $\gamma_{ij}^1 \subset C_j$;
2. any two curves γ_{ij} and $\gamma_{i'j'}$ have only finitely many points in common.

The curves γ_{ij} form a so-called *topological graph* G^* on the vertex set $\{p_1, \dots, p_n\}$, that is, a graph drawn in the plane with curvilinear edges that may cross one another. Clearly, G^* is a “drawing,” a particular embedding of the “abstract” graph $G = G(\mathcal{C})$. See Figure 1 for an illustration.

We need the following result of Pach, Radoičić, and Tóth [8].

Lemma 2.1 (Pach et al. [8]) *Let G^* be a topological graph with n vertices, containing no m pairwise crossing edges with distinct endpoints. Then the number of edges of G^* satisfies*

$$|E(G^*)| \leq c_m n \log^{2(m-3)} n,$$

where c_m is a suitable constant. \square

Suppose that $|E(G^*)| = |E(G)| > c_m n \log^{2(m-3)} n$, where $m = m(k) > 1$ is an integer to be specified later. By Lemma 2.1, we obtain that there are distinct indices $i(1) < j(1), \dots, i(m) < j(m)$ such that the arcs $\gamma_{i(1)j(1)}, \dots, \gamma_{i(m)j(m)} \in E(G^*)$ are pairwise crossing.

For any pair of distinct arcs $\gamma_{i(s)j(s)}, \gamma_{i(t)j(t)}$, at least one of the following four relations holds:

$$\begin{aligned} \gamma_{i(s)j(s)}^0 \cap \gamma_{i(t)j(t)}^0 &\neq \emptyset, \\ \gamma_{i(s)j(s)}^0 \cap \gamma_{i(t)j(t)}^1 &\neq \emptyset, \\ \gamma_{i(s)j(s)}^1 \cap \gamma_{i(t)j(t)}^0 &\neq \emptyset, \\ \gamma_{i(s)j(s)}^1 \cap \gamma_{i(t)j(t)}^1 &\neq \emptyset. \end{aligned}$$

By changing the labeling of the sets C_i and, consequently, of the points $p_i \in V(G^*)$, if necessary, we can assume without loss of generality that for at least one quarter of the pairs $s < t$, we have

$$\gamma_{i(s)j(s)}^0 \cap \gamma_{i(t)j(t)}^0 \neq \emptyset.$$

Construct a bipartite graph B on the vertex set $V(B) = \{1, 2, \dots, m\} \cup \{1', 2', \dots, m'\}$ by connecting s and t' with an edge if and only if $s < t$ and

$$\gamma_{i(s)j(s)}^0 \cap \gamma_{i(t)j(t)}^0 \neq \emptyset.$$

We have

$$|E(B)| \geq \frac{1}{4} \binom{m}{2}.$$

Applying the Kővári–Sós–Turán theorem quoted in the Introduction, we obtain that B has a complete bipartite subgraph $K_{k,k}$, provided that

$$\frac{1}{4} \binom{m}{2} > b_k (2m)^{2-\frac{1}{k}}.$$

Choose $m = m(k) = O(b_k)^k$ to be the smallest positive integer that satisfies the last inequality. We can conclude that there exist vertices $s_1, \dots, s_k, t'_1, \dots, t'_k \in V(B)$ with the property that every element of

$$\{\gamma_{i(s_1)j(s_1)}^0, \dots, \gamma_{i(s_k)j(s_k)}^0\}$$

intersects all elements of

$$\{\gamma_{i(t'_1)j(t'_1)}^0, \dots, \gamma_{i(t'_k)j(t'_k)}^0\}.$$

It follows from our construction that $\gamma_{i(r)j(r)}^0 \subset C_{i(r)}$, for every $1 \leq r \leq m$. Therefore, the corresponding collections of sets $\{C_{i(s_1)}, \dots, C_{i(s_k)}\}$ and $\{C_{i(t'_1)}, \dots, C_{i(t'_k)}\}$ induce a complete bipartite graph $K_{k,k}$ in the intersection graph $G(\mathcal{C})$, which leads to a contradiction that completes the proof of Theorem 1. \square

3 Intersection graphs of convex sets: Proof of Theorem 2

The polylogarithmic factor in the bound of Theorem 1 is a consequence of the fact that our proof was based on Lemma 2.1, a general statement on *topological* graphs. If the elements of \mathcal{C} can be arbitrary connected sets, the particular drawing of $G(\mathcal{C})$ we constructed may be quite complicated. However, if \mathcal{C} consists of *convex* sets, one can explore some simple structural properties of this drawing. Specifically, it will be sufficient to consider *x-monotone* topological graphs, that is, topological graphs with the property that any vertical line (parallel to the y -axis) intersects every edge at most once. In this special case, Valtr [9] managed to reduce the exponent of the polylogarithmic factor in Lemma 2.1 to one.

Lemma 3.1 (Valtr [9]) *Let G^* be an x -monotone topological graph with n vertices, containing no m pairwise crossing edges with distinct endpoints. Then the number of edges of G^* satisfies*

$$|E(G^*)| \leq c'_m n \log n,$$

where c'_m is a suitable constant. \square

Suppose that \mathcal{C} is a collection of compact convex sets in the plane. We adapt the proof of Theorem 1 to this case. First we show that a slightly modified version of the graph G^* we used in the proof of Theorem 1 can be drawn as an x -monotone topological graph.




Figure 2: The drawing of G^* in the convex case.

For any $C \in \mathcal{C}$, let l_C, r_C denote the leftmost and rightmost points of C . Without loss of generality, we may assume that these points are unique, and that the set $P = \{l_C, r_C \mid C \in \mathcal{C}\}$ consists of $2n$ distinct points, no three of which are collinear. We now define the topological graph G^* on the vertex set P as follows. For each pair of intersecting sets $C_i, C_j \in \mathcal{C}$ with $i < j$, connect l_{C_i} and r_{C_j} by a polygonal path γ_{ij} consisting of two segments that meet at an arbitrary point of $C_i \cap C_j$. Clearly, γ_{ij} is an x -monotone curve, and we have $|E(G^*)| = |E(G)|$. See Figure 2.

Suppose that $|E(G^*)| = |E(G)| > c'_m \cdot 2n \log(2n)$, where $m = m(2k) > 1$ denotes the same integer as in the proof of Theorem 1, but defined for $2k$ rather than for k . By Lemma 3.1, we obtain that there are distinct indices $i(1) < j(1), \dots, i(m) < j(m)$ such that the arcs $\gamma_{i(1)j(1)}, \dots, \gamma_{i(m)j(m)} \in E(G^*)$ are pairwise crossing. The proof now proceeds in essentially the same way as the proof of Theorem 1. The only difference is that at the very end of the proof, when we find a complete bipartite graph $K_{2k,2k}$ in the intersection graph $G(\mathcal{C})$, some of the vertices in one color class of $K_{2k,2k}$ may represent the same member $C \in \mathcal{C}$ as a vertex in the other color class (because each C is represented in G^* by two vertices). However, we can always choose a complete bipartite subgraph $K_{k,k} \subset K_{2k,2k}$, whose vertices represent distinct sets. This leads to a contradiction that establishes Theorem 2. \square

4 Forbidden C_4 : Proofs of Theorem 3 and Corollary 4

First we establish Theorem 3. Let \mathcal{C} be a collection of n convex sets in the plane such that their intersection graph $G(\mathcal{C})$ contains no C_4 .

For any $C \in \mathcal{C}$, let s_C denote the straight-line segment connecting the leftmost point and the rightmost point of C . We refer to s_C as the *spine* of C . Let \mathcal{S} be the set of the spines of all the sets in \mathcal{C} . By slightly perturbing the sets if necessary, without changing the intersection graph $G(\mathcal{C})$, we can achieve that the leftmost and rightmost points of the sets C are unique and that the spines are in general position, that is, no three of their endpoints are collinear and no three of them pass through the same point.

Let Ξ denote the *vertical decomposition* of $\mathcal{A}(\mathcal{S})$, obtained by erecting a vertical segment up and down from each endpoint of a spine and from each intersection point of two spines, and extending these segments until they hit another spine, or else all the way to infinity




Figure 3: The drawing of G^* in the convex case.

[2]. Each cell Δ of Ξ is a vertical trapezoid, bounded by portions of spines on the top and bottom, and by vertical segments to the left and to the right; any of these boundary segments may be missing. Let X denote the number of intersection points of the spines. By assumption, all intersection points are simple, therefore, it is easy to verify that the number of trapezoids (cells) in Ξ is equal to $3n + 3X + 1 = O(n + X)$.

Fix a cell Δ , and let $A, B \in \mathcal{C}$ be the sets such that the top (resp., bottom) boundary of Δ is contained in s_A (resp., s_B). Let $K \in \mathcal{C}$ be an object that intersects Δ , and let p be a point in $K \cap \Delta$. The vertical line λ through p must intersect the spine s_K of K . Suppose that $K \neq A, B$ and that s_K passes below s_A within the vertical slab spanned by Δ .

Suppose that λ intersects the spine s_C of another set $C \in \mathcal{C}$ between s_A and s_K . Among all such sets, we fix C to be the set whose spine is nearest to s_A along λ . Then A and C uniquely determine K . That is, there cannot exist another $K' \in \mathcal{C}$, $K' \neq A, C, K$, such that $s_{K'}$ lies below s_A and the vertical line through a point in $K' \cap \Delta$ meets s_C between s_A and $s_{K'}$. This follows by noting that such a K' would have formed, together with A, C , and K , a forbidden C_4 in $G(\mathcal{C})$. Since s_C is the spine directly preceding s_A along λ , s_A and s_C form a cell Δ' of Ξ that lies below Δ , is crossed by λ , and is a neighbor of Δ in the planar map Ξ . See Figure 3 for an illustration.

A symmetric argument shows that if $L \in \mathcal{C}$ is an object that intersects Δ at some point q , with s_L lying above s_B within the vertical slab spanned by Δ , then either s_L lies directly above s_B , or else B and the set $D \in \mathcal{C}$, whose spine lies directly above s_B along the vertical line through q , uniquely determine L .

Sweep Δ by a vertical line λ from left to right. At any step during the sweep, λ passes through a bottom neighbor cell Δ' and through a top neighbor cell Δ'' . The triple $(\Delta, \Delta', \Delta'')$ changes only when either the top neighbor or the bottom neighbor changes, implying that the number of such triples, with Δ fixed, is at most the number of top and bottom neighbors of Δ in Ξ . Since Ξ is a planar map whose dual map has no digonal faces, we conclude that the overall number of such triples is $O(|\Xi|) = O(n + X)$.

As argued above, a fixed triple $(\Delta, \Delta', \Delta'')$ determines at most six sets of \mathcal{C} that can intersect Δ at points p that lie “in between” Δ' and Δ'' , in the sense that the vertical line through p intersects both Δ' and Δ'' . Hence, the number of pairs of sets in \mathcal{C} that can




Figure 4: Two triangular faces cannot share a vertex.

intersect within this portion of Δ is $O(1)$. This implies that the number of edges of $G(\mathcal{C})$ is $O(n + X)$.

To complete the proof of Theorem 3, it remains to show that $X = O(n)$.

Lemma 4.1 *Let \mathcal{S} be a set of n straight-line segments in general position in the plane, whose intersection graph has no C_4 as a subgraph. Then the number of crossings between the elements of \mathcal{S} satisfies $X \leq 21n$.*

Proof: Consider the arrangement $\mathcal{A}(\mathcal{S})$ of \mathcal{S} as a planar map M , that is, introduce a new vertex at each crossing. Clearly, M has no *digons* (faces with two sides) and no quadrilateral faces, for this would correspond to a C_4 in $G(\mathcal{C})$. The number of vertices of M is $V = 2n + X$, and the number of its edges is $E = n + 2X$. Let F and F_3 denote the number of faces and the number of triangular faces of M , respectively. Obviously, we have $3F_3 + 5(F - F_3) \leq 2E$, or, equivalently,

$$F \leq \frac{2}{5}E + \frac{2}{5}F_3 \leq \frac{2}{5}(n + 2X) + \frac{2}{5}F_3.$$

Applying Euler's formula, we obtain

$$2 \leq V - E + F \leq (2n + X) - (n + 2X) + \frac{2}{5}(n + 2X) + \frac{2}{5}F_3,$$

which implies that

$$X \leq 7n + 2F_3 - 10. \tag{1}$$

It remains to bound F_3 . Since we assumed that the segments were in general position, no two triangular faces f, f' can share a vertex w , unless they lie in opposite wedges formed by the two spines meeting at w . However, this case cannot arise, because then the four spines that bound f and f' would induce a forbidden C_4 in $G(\mathcal{C})$; see Figure 4. Hence, the F_3 triangular faces define $3F_3$ *distinct* crossing points, so $3F_3 \leq X$. Substituting this in (1) yields

$$X \leq 7n + 2F_3 - 10 \leq 7n + \frac{2}{3}X - 10, \quad \text{or} \quad X \leq 21n - 30.$$

This completes the proof of Lemma 4.1 and hence of Theorem 3. \square

To establish Corollary 4, we have to recall the so-called Crossing Lemma of Ajtai et al. [1] and Leighton [6].

Lemma 4.2 *Any topological graph with n vertices and $e > 4n$ edges determines at least $\frac{e^3}{64n^2}$ crossings. \square*

Consider now a geometric graph G with n vertices and e edges, determining no 2×2 grid. Notice that the intersection graph of the set of segments obtained from the edges of G by clipping them near their endpoints has no C_4 . Thus, we can apply Lemma 4.1 to obtain that the number of crossings between the edges satisfies $X \leq 21e$. On the other hand, Lemma 4.2 implies that $X \geq \frac{e^3}{64n^2}$, provided that $e > 4n$. Comparing these two bounds, we conclude that $e < 40n$, completing the proof of Corollary 4. \square

References

- [1] M. Ajtai, V. Chvátal, M. Newborn and E. Szemerédi, Crossing-free subgraphs, *Ann. Discrete Math.* **12** (1982), 9–12.
- [2] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf, *Computational Geometry: Algorithms and Applications (2nd ed.)*, Springer-Verlag, Berlin, 2000.
- [3] A. S. Besicovitch, On Crum’s problem, *J. London Math. Soc.* **22** (1947), 285–287.
- [4] P. Erdős and A. H. Stone, On the structure of linear graphs, *Bull. Amer. Math. Soc.* **52** (1946), 1087–1091.
- [5] T. Kővári, V. T. Sós, and P. Turán, On a problem of K. Zarankiewicz, *Colloquium Math.* **3** (1954), 50–57.
- [6] F. T. Leighton, *Complexity Issues in VLSI*, MIT Press, Cambridge, MA, 1983.
- [7] J. Pach, R. Pinchasi, M. Sharir, and G. Tóth, Topological graphs with no large grids, *Graphs and Combinatorics* **21** (2005), 355–364.
- [8] J. Pach, R. Radoičić, and G. Tóth, Relaxing planarity for topological graphs, *Discrete and Computational Geometry, Lecture Notes in Comput. Sci.*, **2866**, Springer-Verlag, Berlin, 2003, 221–232.
- [9] P. Valtr, Graph drawing with no k pairwise crossing edges, in: G. DiBattista (ed.) *Graph Drawing ‘97, Lecture Notes in Comput. Sci.* **1353**, Springer-Verlag, Berlin, 1997, 205–218.