On the number of touching pairs in a set of planar curves

Péter Györgyi[a,b], Bálint Hujter[a], Sándor Kisfaludi-Bak[c,*]

[a] *Institute of Mathematics, Eötvös University, Budapest, Hungary*
[b] *Institute for Computer Science and Control, Budapest, Hungary*
[c] *Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands*

Abstract

Given a set of planar curves (Jordan arcs), each pair of which meets – either crosses or touches – exactly once, we establish an upper bound on the number of touchings. We show that such a curve family has $O(t^2 n)$ touchings, where t is the number of faces in the curve arrangement that contains at least one endpoint of one of the curves. Our method relies on finding special subsets of curves called quasi-grids in curve families; this gives some structural insight into curve families with a high number of touchings.

Keywords: Combinatorial geometry, Touching curves, Pseudo-segments

1. Introduction

The combinatorial examination of incidences in the plane has proven to be a fruitful area of research. The first seminal results are the crossing lemma that establishes a lower bound on the number of edge crossings in a planar drawing of a graph (Ajtai et al., Leighton [1, 2]), and the theorem by Szemerédi and Trotter [3], concerning the number of incidences between lines and points. Soon, the incidences of more general geometric objects (segments, circles, algebraic curves, pseudo-circles, Jordan arcs, etc.) became the center of attention [4, 5, 6, 7, 8, 9]. With the addition of curves, the distinction between touchings and crossings is in order.

Usually, the curves are either Jordan arcs, i.e., the image of an injective continuous function $\varphi : [0, 1] \to \mathbb{R}^2$, or closed Jordan curves, where φ is injective on $[0, 1)$ and $\varphi(0) = \varphi(1)$. Generally, it is supposed that the curves intersect in a finite number of points, and that the curves are in general position: three curves cannot meet at one point, and (in case of Jordan arcs) an endpoint of a curve does not lie on any other curve. (For technical purposes, we will allow curve endpoints to coincide in some proofs.)

*Corresponding author

Email addresses: gyorgyip@cs.elte.hu (Péter Györgyi), hujterb@cs.elte.hu (Bálint Hujter), s.kisfaludi.bak@tue.nl (Sándor Kisfaludi-Bak)

Let P be a point where curve a and b meet. Take a circle γ with center P and a small enough radius so that it intersects both a and b twice, and the disk determined by γ is disjoint from all the other curves , and contains no other intersections of a and b. Label the intersection points of γ and the two curves with the name of the curve. We say that a and b *cross* in P if the cyclical permutation of labels around γ is *abab*, and a and b *touch* in P if the cyclical permutation of labels is *aabb*. In a family of curves, let X be the set of crossings and T be the set of touchings.

The Richter-Thomassen conjecture [10] states that given a collection of n pairwise intersecting closed Jordan curves in general position in the plane, the number of crossings is at least $(1 - o(1))n^2$. A proof of the Richter-Thomassen conjecture has recently been published by Pach et al. [11]. They show that the same result holds for Jordan arcs as well.

It would be preferable to get more accurate bounds for the ratio of touchings and crossings. Fox et al. constructed a family of x-monotone curves with ratio $|X|/|T| = O(\log n)$ [12]. If we restrict the number of intersections between any two curves, then it is conjectured that the above ratio is much higher. It has been shown that a family of intersecting pseudo-circles (i.e., a set of closed Jordan-curves, any two of which intersect exactly once or twice) has at most $O(n)$ touchings [7]. We would like to examine a similar statement for Jordan arcs.

A family of Jordan arcs in which any pair of curves intersect at most once (apart from the endpoints) will be called a *family of pseudo-segments.* Our starting point is this conjecture of János Pach [13]:

Conjecture 1. Let $\mathcal{C}$ be a family of pseudo-segments. Suppose that any pair of curves in $\mathcal{C}$ intersect exactly once. Then the number of touchings in $\mathcal{C}$ is $O(n)$.

A family of pseudo-segments is *intersecting* if every pair of curves intersects (i.e., either touches or crosses) exactly once outside their endpoints.

Two important special cases of the above are the cases of *grounded* and *double-grounded* curves. (The definitions are taken verbatim from [9].) A collection $\mathcal{C}$ of curves is *grounded* if there is a closed Jordan curve g called *ground* such that each curve in $\mathcal{C}$ has one endpoint on g and the rest of the curve is in the exterior of g. The collection is *double grounded* if there are disjoint closed Jordan curves g_1 and g_2 such that each curve $c \in \mathcal{C}$ has one endpoint on g_1 and the other endpoint on g_2, and the rest of c is disjoint from both g_1 and g_2.

According to our knowledge the best upper bound is $O(n \log n)$ for the number of touchings in a double-grounded x-monotone family of pseudo-segments [14] and we do not know any (non-trivial) result for the grounded case.

1.1. Our contribution

Let $\mathcal{C}$ be an intersecting family of pseudo-segments. There is a planar graph drawing that corresponds to this family: the vertices are the crossings and touchings, and the edges are the sections of the curves between neighboring intersections. (Notice that the sections between curve endpoints and the neighboring intersections are not represented in this graph.) Consider the faces of this planar graph drawing. Let $t_{\mathcal{C}}$ be the number of faces that contain an endpoint of at least one curve in $\mathcal{C}$. Our main theorem can be stated as follows:

Theorem 2. *Let $\mathcal{C}$ be an n-element intersecting family of pseudo-segments on the Euclidean plane. Then the number of touchings between the curves is $f(n) = O(t_{\mathcal{C}}^2 n)$.*

If $t_{\mathcal{C}}$ is constant, this theorem settles Conjecture 1. Note that this includes the case when $\mathcal{C}$ is a double-grounded intersecting family of pseudo-segments:

Corollary 3. *Let $\mathcal{C}$ be an n-element double-grounded intersecting family of pseudo-segments. Then the number of touchings between the curves is $O(n)$.*

A careful look at the proof of the main theorem yields the following result for grounded intersecting families of pseudo-segments:

Theorem 4. *Let $\mathcal{C}$ be an n-element grounded intersecting family of pseudo-segments. Then the number of touchings between the curves is $O(t_{\mathcal{C}} n)$.*

The intuition behind our approach can be described as follows. Curves starting in the same face of an arrangement can be thought of as curves having the same endpoints. A curve going from point A to B that touches some other curve g can do that touching only in a constant number of ways, depending on which side of g is touched and in which direction. We observe that a collection of curves going from A to B must therefore contain a subcollection that touch g the same way, and these curves must have a very special grid-like structure, which we call *quasi-grids*.

It turns out that quasi-grids always emerge when we take two grid families of pseudo-segments, one containing curves from A to B, the other containing curves from C to D. Note that a curve touching all curves in a large quasi-grid has to lie outside the "grid cells", since it cannot cross the quasi-grid curves, and within a "grid cell" it could only reach at most four curves. If we find two curves touching the same large quasi-grid, then (intuitively) those two curves would have many intersections – this is not possible in an intersecting family of pseudo-segments. We show that the number of touchings between a pair of fixed endpoint curve families is linear in the size of these families. We then use this observation to get the bound on the total number of touchings.

2. Proof of the main theorem

The rigorous proof of our main theorem is based upon a key lemma. Its proof anticipates and uses several technical lemmas which are detailed in Sections 3 and 4.

Before stating the key lemma, we introduce some notations. The notation $g \asymp h$ means that curves g and h touch each other. If A and B are (not necessarily distinct) points on the plane, then $\mathcal{C}(A, B)$ denotes the set of directed curves going from A to B. Note that here we consider curves as directed ones for technical reasons (for example, we can refer to the sides of a directed curve as *left* and *right*).

Lemma 5. *Let A, B, C, D be not necessarily distinct points on the plane, and $\mathcal{C}_1$ and $\mathcal{C}_2$ be finite disjoint curve families from $\mathcal{C}(A, B)$ and $\mathcal{C}(C, D)$, respectively. If $\mathcal{C}_1 \cup \mathcal{C}_2$ is an intersecting family of pseudo-segments, then*

1. *the number of $c_1 \asymp c_2$ touchings where $c_1 \in \mathcal{C}_1$ and $c_2 \in \mathcal{C}_2$ is $O(|\mathcal{C}_1 \cup \mathcal{C}_2|)$;*
2. *the number of touchings between curves of $\mathcal{C}_i$ is $O(|\mathcal{C}_i|)$ $(i = 1, 2)$.*

Proof. We only consider the first claim, the second can be proven with the same tools. Suppose for contradiction that there are $\omega(|\mathcal{C}_1 \cup \mathcal{C}_2|)$ instances of $c_1 \asymp c_2$ touchings.

Let K be a large constant. Without loss of generality, we can suppose that each curve of $\mathcal{C}_i$ touches at least K curves of $\mathcal{C}_j$. To see this, consider first the bipartite graph G with vertex set $\mathcal{C}_1 \cup \mathcal{C}_2$, where the edges correspond to the $c_1 \asymp c_2$ touchings ($c_1 \in \mathcal{C}_1$ and $c_2 \in \mathcal{C}_2$). If G has vertices of degree less than K, then delete those vertices and the incident edges. Iterate this procedure until the minimum degree is at least K or the graph is empty. If G had at least $K|\mathcal{C}_1 \cup \mathcal{C}_2|$ edges, then this procedure cannot result in an empty graph.

Let $g \in \mathcal{C}_1$ be an arbitrary curve. By Lemma 10, there is a *quasi-grid* with respect to g formed by at least $K/48 > 3$ curves. A quasi-grid is depicted in Figure 1, the precise definition is given in Definition 6.

Consider an "inner" curve h in this quasi-grid. By Lemma 11, if a curve touches h, then it must also touch g or a neighboring curve of h in the quasi-grid. By our starting assumption, at least K curves touch h. Then by Lemma 10, at least $K/48$ of the curves touching h must also touch another specific curve h', and at least $K/(48)^2$ of these form a quasi-grid $\mathcal{Q}$ with respect to both h and h'.

Therefore, by choosing $K \geq 4 \cdot 48^2 + 1$, the quasi-grid $\mathcal{Q}$ can be forced to contain at least five curves. This is a contradiction by Lemma 13. □

Next we show how Lemma 5 implies Theorem 2. Let $t = t_{\mathcal{C}}$.

Proof of Theorem 2. Consider the planar graph drawing that corresponds to $\mathcal{C}$. Let the faces of this planar graph drawing that contain an endpoint of at least one curve in $\mathcal{C}$ be: $F_1, F_2, \ldots, F_t$.

For $i = 1, 2, \ldots, t$, let P_i be an arbitrary point in the interior of F_i not incident to any curve in $\mathcal{C}$. Each curve endpoint inside F_i can be connected to P_i without adding any intersections between the curves of $\mathcal{C}$ with the exception of P_i. Let $\mathcal{C}'$ be the family of pseudo-segments obtained from $\mathcal{C}$ by this procedure.

Partition $\mathcal{C}'$ to disjoint subsets $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_s$ so that two curves are in the same subset if and only if their endpoints are the same. Note that $s \leq \binom{t+1}{2}$. Fix the orientation of each curve in $\mathcal{C}'$ from P_i to P_j if $i < j$, and arbitrarily if $i = j$.

Let f_k denote the number of touchings inside $\mathcal{C}_k$ and $f_{k,l}$ denote the number of touchings between $\mathcal{C}_k$ and $\mathcal{C}_l$. Then the total number of touchings in $\mathcal{C}'$ is

$$f(n) = \sum_k f_k + \sum_{k<l} f_{k,l} = \sum_k O(|\mathcal{C}_k|) + \sum_{k<l} O(|\mathcal{C}_k| + |\mathcal{C}_l|)$$
$$= O(n) + \sum_k (s-1)O(|\mathcal{C}_k|) = O(sn) = O\left(t^2 n\right),$$

where the second equation follows from Lemma 5. □

Notice that in case of a grounded intersecting family of pseudo-segments, we have $s = t + 1$, so $O(sn) = O(tn)$, which proves Theorem 4.

3. Quasi-grids and their occurrence

3.1. Notations and definitions

We introduce several notations used in the paper. Let g and h be a pair of directed curves. If g touches the left side of h, and they have the same direction at the touching point, then write $g \upuparrows h$. (More precisely, let γ be a circle around the intersection P with a small enough radius so that it intersects both a and b twice, and the disk determined by γ is disjoint from all the other curves, and contains no other intersections of a and b. We label the points where g and h enters γ by g and h, and assign the labels g' and h' to the points where they exit. We say that the right side of g touches the left side of h in P if the counter-clockwise cyclic order of labels on γ is $ghh'g'$.) Notice that this relation is not symmetric, i.e., $g \upuparrows h \not\Leftrightarrow h \upuparrows g$. If g and h have different directions at the touching point (so the counter-clockwise cyclic order of labels on γ is $gg'hh'$ or $gh'hg'$), then write $g \downharpoonleft\upharpoonright h$ or $g \upharpoonleft\downharpoonright h$ depending on which side of h is touched by g. We say that c_1 and c_2 are *g-touch equivalent* if they touch g on the same side and in the same direction, i.e., $(g \upuparrows c_1 \wedge g \upuparrows c_2)$ or $(g \downharpoonleft\upharpoonright c_1 \wedge g \downharpoonleft\upharpoonright c_2)$ or $(c_1 \upuparrows g \wedge c_2 \upuparrows g)$ or $(c_1 \upharpoonleft\downharpoonright g \wedge c_2 \upharpoonleft\downharpoonright g)$. A set of curves is g-touch equivalent if its elements are pairwise g-touch-equivalent.

For a directed curve g with points A and B that lie on the curve in this order, let $A \xrightarrow{g} B$ be the closed directed subcurve from A to B, and $B \xleftarrow{g} A$ will denote the reverse directed subcurve from B to A. This notation can be iterated, e.g. if $P \in h \cap g$, then $A \xrightarrow{g} P \xleftarrow{h} Q$ denotes the curve which starts from $A \in g$, continues on g to the intersection point P, then changes to h, and goes on h in reverse direction until it ends in $Q \in h$. When referring to undirected subcurves, we use $A \overset{g}{\text{---}} B$. Sometimes these notations are also used to denote the ordering of points on a particular curve.

As already defined, $\mathcal{C}(A, B)$ is the set of directed curves going from A to B. For a curve $c \in \mathcal{C}(A, B)$, let $c^\star = c \setminus \{A, B\}$. For a set of curves $\mathcal{C} = \{c_1, c_2, \dots c_k\}$, let $\mathcal{C}^\star = \{c_1^\star, c_2^\star, \dots c_k^\star\}$.

The objects called quasi-grids are the main tool of this paper. Intuitively, the below definition says that the incidences of a quasi-grid are exactly as shown in Figure 1, with the exception of the points X, Y, A and B — we allow these to coincide arbitrarily.

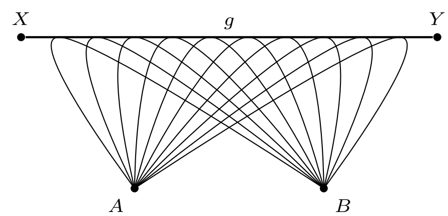

Figure 1: A quasi-grid for the case $g \upuparrows c_i$. Swapping X, Y or A, B gives the other 3 cases.

Definition 6 (Quasi-grid)**.** A set of curves $\mathcal{C} = \{c_1, c_2, \ldots, c_k\} \subseteq \mathcal{C}(A, B)$ forms a *quasi-grid* with respect to a curve $g \in \mathcal{C}(X, Y)$ if:

1. $\mathcal{C}^\star \cup \{g^\star\}$ is an intersecting family of pseudo-segments
2. $\mathcal{C}^\star$ is g-touch-equivalent with touching points $P_i = g \cap c_i$
3. $P_{i,j} = c_i^\star \cap c_j^\star$ is a crossing point
4. the ordering of points on g is $P_1 \overset{g}{\text{---}} P_2 \overset{g}{\text{---}} \ldots \overset{g}{\text{---}} P_k$
5. the ordering of points on c_j $(j = 1, 2, \ldots, k)$ is

$$A \xrightarrow{c_j} P_{1,j} \xrightarrow{c_j} P_{2,j} \xrightarrow{c_j} \ldots \xrightarrow{c_j} P_{j-1,j} \xrightarrow{c_j} P_j \xrightarrow{c_j} P_{j,j+1} \xrightarrow{c_j} \ldots \xrightarrow{c_j} P_{j,k} \xrightarrow{c_j} B.$$

An example for a quasi-grid can be seen in Figure 1. Throughout the paper (if we do not indicate it otherwise) we assume that the indices of the curves in $\mathcal{C}$ describe the order of their touching points on g.

3.2. Finding quasi-grids in curve configurations

The goal of this subsection is to prove that in a family of pseudo-segments, the set of g-touch equivalent curves with given endpoints form a constant number of quasi-grids. Intuitively, Lemma 7 shows that in a family of pseudo-segments, the g-touch equivalent curves from $\mathcal{C}(A, B)$ (where $A \neq B$) can still have two distinct types. Note that these types cannot be defined separately, only in relation to each other. In Lemma 8, we establish that the curves in each type form a quasi-grid with respect to g. Lemma 9 examines the case $A = B$.

Lemma 7. *Fix a curve $g \in \mathcal{C}(X, Y)$ and suppose that c_1 and c_2 are g-touch-equivalent curves from $\mathcal{C}(A, B)$ with touching points P_1 and P_2 respectively, where $A \neq B$. Note that P_1 and P_2 divide c_1 and c_2 into their first and second parts. Suppose further that $\{c_1^\star, c_2^\star, g^\star\}$ is a family of pseudo-segments. Then $c_1^\star$ crosses $c_2^\star$ at a point Q, which is either the intersection of the first part of c_1 with the second part of c_2, or vice versa: the intersection of the second part of c_1 with the first part of c_2.*

Proof. Suppose (without loss of generality) that $g \upuparrows c_1$, $g \upuparrows c_2$, and that P_1 precedes P_2 on g. Consider the closed directed curve $\ell = A \xrightarrow{c_1} P_1 \xrightarrow{g} P_2 \xleftarrow{c_2} A$ (curves with gray halo in the middle and right part of Figure 2). We show that

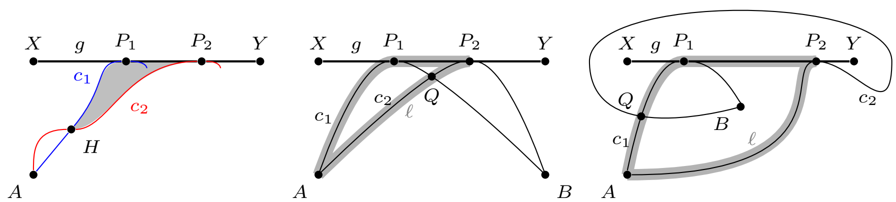

Figure 2: Left: c_1 and c_2 cannot intersect before reaching g; middle and right: the possible configurations

ℓ is a Jordan-curve. Suppose for contradiction that $A\xrightarrow{c_1}P_1$ and $A\xrightarrow{c_2}P_2$ has an intersection point $H \neq A$ (see the left picture in Figure 2). Since $\{c_1^\star, c_2^\star, g^\star\}$ is a family of pseudo-segments, there can be no further intersections between $c_1^\star$ and $c_2^\star$. It follows that $\ell' = H\xrightarrow{c_1}P_1\xrightarrow{g}P_2\xleftarrow{c_2}H$ is a Jordan curve that separates the plane into its left and right (shaded) side regions. Notice that $P_1\xrightarrow{c_1}B$ begins in the right side region of ℓ', while $P_2\xrightarrow{c_2}B$ begins in the left side region by our assumptions $g \upuparrows c_1$ and $g \upuparrows c_2$. Since $P_1\xrightarrow{c_1}B\xleftarrow{c_2}P_2$ is a continuous curve that begins and ends in different sides of ℓ', it must cross ℓ' in a point distinct from H; we arrived at a contradiction.

Thus $(A\xrightarrow{c_1}P_1) \cap (A\xrightarrow{c_2}P_2) = \{A\}$, hence ℓ is a Jordan-curve. By similar argument as above, $P_1\xrightarrow{c_1}B\xleftarrow{c_2}P_2$ is a continuous curve that begins and ends in different sides of ℓ, so it must cross ℓ. Since c_1, c_2 and g are not self-intersecting and P_1 and P_2 already account for the intersections between g and $\{c_1 \cup c_2\}$, the only remaining possibilities are that the crossing point is as claimed, i.e., $Q = (A\xrightarrow{c_2}P_2) \cap (P_1\xrightarrow{c_1}B)$ or $Q = (A\xrightarrow{c_1}P_1) \cap (P_2\xrightarrow{c_2}B)$ (see the middle and the right picture in Figure 2). ☐

Notice that the above lemma states that the curve c_2 meets c_1 before it meets g if and only if the first part of c_2 crosses the second part of c_1 and vice versa: c_2 meets g before it meets c_1 if and only if the second part of c_2 crosses the first part of c_1. This equivalence will be used several times in the following lemmas.

Lemma 8. *Let $g \in \mathcal{C}(X, Y)$, and let $\mathcal{H}$ be a set of g-touch-equivalent curves from $\mathcal{C}(A, B)$, where $A \neq B$. If $\mathcal{H}^\star \cup \{g^\star\}$ is a family of pseudo-segments, then $\mathcal{H}$ is the disjoint union of at most two quasi-grids with respect to g.*

Proof. We deal with the case $g \upuparrows h$ for all $h \in \mathcal{H}$, the other three cases are similar. Let $h \in \mathcal{H}$ be the curve that has the first touching point on $X\xrightarrow{g}Y$ among the curves from $\mathcal{H}$. Let $\mathcal{H}_1 \subseteq \mathcal{H}$ consist of h and the curves from $\mathcal{H}$ that meet h before they meet g. Let $\mathcal{H}_2 := \mathcal{H} \setminus \mathcal{H}_1$. We prove that $\mathcal{H}_1$ and $\mathcal{H}_2$ are both quasi-grids with respect to g.

Let $\mathcal{H}_1 = \{h = h_1, h_2, \dots, h_\ell\}$ and let $P_i = g \cap h_i$. Assume without loss of generality that $P_1\xrightarrow{g}P_2\xrightarrow{g}\dots\xrightarrow{g}P_\ell$. We show that $\mathcal{H}_1$ is a quasi-grid with respect to g. (See Figure 3.)

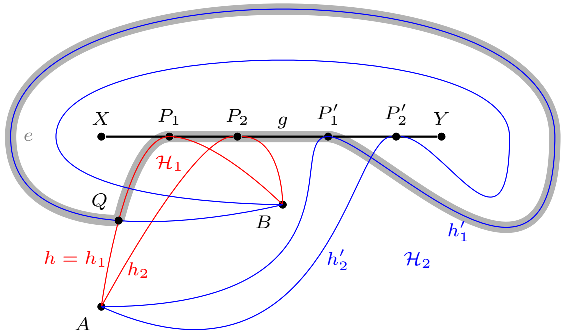

Figure 3: Two quasi-grids with respect to g.

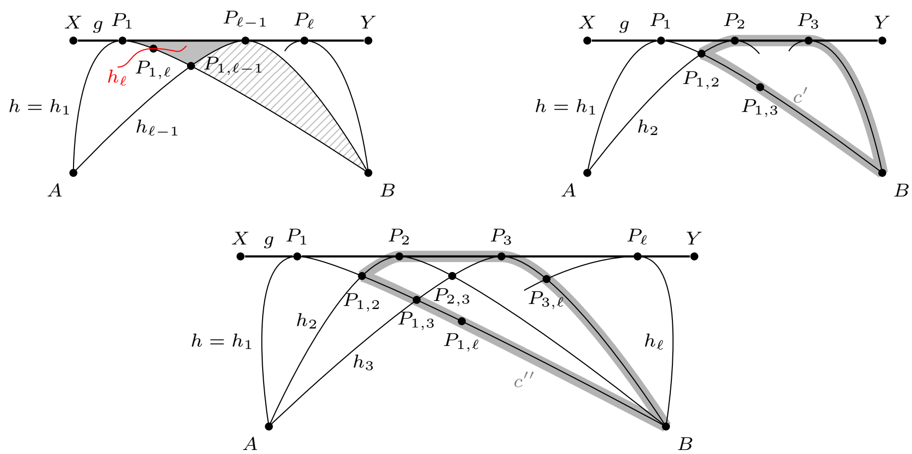

Figure 4: Top left: if $P_{1,\ell}$ is on $P_1 \xrightarrow{h_1} P_{1,\ell-1}$; top right: the case $|\mathcal{H}_1| = \ell = 3$; bottom: the case $\ell \geq 4$.

The proof is by induction on the number of curves in $\mathcal{H}_1$, for $\ell = 1$ the statement is trivial. Lemma 7 yields the statement for $\ell = 2$.

We claim that h_ℓ crosses h_1 between $P_{1,\ell-1}$ and B. By the definition of $\mathcal{H}_1$ and Lemma 7, $P_{1,\ell}$ must lie on $P_1 \xrightarrow{h_1} B$. Suppose for contradiction that the ordering on h_1 is $P_1 \xrightarrow{h_1} P_{1,\ell} \xrightarrow{h_1} P_{1,\ell-1}$ (top left of Figure 4). Consider the closed Jordan curve $c = P_1 \xrightarrow{g} P_{\ell-1} \xleftarrow{h_{\ell-1}} P_{1,\ell-1} \xleftarrow{h_1} P_1$. (The right side region of c is shaded.)

Notice that A and B are on the left side of c. To see this, consider that c is made up of three curve segments, and there can be no further intersections among these three curves, so h_1 and $c \setminus (P_1 \xrightarrow{h_1} P_{1,\ell-1})$ are disjoint. Since the type of touching at P_1 is $g \upuparrows h_1$, we can see that $(A \xrightarrow{h_1} P_1) \setminus P_1$ and consequently point A in particular lies in the left side region of c. A similar argument for $h_{\ell-1}$ shows that B is also in the left side region.

Since h_ℓ is already crossing c once at $P_{1,\ell}$, it has to cross it one more time, because its endpoints A and B are on the same side of c. Since h_ℓ touches g outside c and it already has an intersection with h_1, it will cross $P_{\ell-1} \xleftarrow{h_{\ell-1}} P_{1,\ell-1}$. Now h_ℓ has entered the right side of $P_{1,\ell-1} \xrightarrow{h_{\ell-1}} B \xleftarrow{h_1} P_{1,\ell-1}$ (the region with the line pattern), while P_ℓ is on the left side (since $g \upuparrows h_{\ell-1}$). So h_ℓ would have to cross h_1 or $h_{\ell-1}$ one more time, which is a contradiction.

If $\ell = 3$, then consider the curve $c' = P_2 \xrightarrow{g} P_3 \xrightarrow{h_3} B \xleftarrow{h_1} P_{1,3} \xleftarrow{h_1} P_{1,2} \xrightarrow{h_2} P_2$ (see the top right of Figure 4). Since h_2 and h_3 touch g in the same direction, h_3 goes from $P_{1,3}$ to P_3 in the same side of c' where h_2 goes from P_2 to B (or a point on $(P_3 \xrightarrow{h_3} B)$). Thus, h_2 and h_3 cross each other at a point $P_{2,3} = (P_{1,3} \xrightarrow{h_3} P_3) \cap (P_2 \xrightarrow{h_2} B)$. By induction, the points on h_1 and h_2 are also in the required order.

For $\ell \geq 4$ the induction is used both for $h_1, h_2, \ldots, h_{\ell-1}$ and $h_1, h_3, h_4, \ldots, h_\ell$. We only need to show that h_2 and h_ℓ cross each other at a point $P_{2,\ell}$ which satisfies our ordering conditions. Indeed, on the right side of the curve

$$c'' = P_2 \xrightarrow{g} P_3 \xrightarrow{h_3} P_{3,\ell} \xrightarrow{h_3} B \xleftarrow{h_1} P_{1,\ell} \xleftarrow{h_1} P_{1,2} \xrightarrow{h_2} P_2,$$

there is a crossing $P_{2,\ell} = (P_{1,\ell} \xrightarrow{h_\ell} P_{3,\ell}) \cap (P_2 \xrightarrow{h_2} B)$: this shows that the ordering on h_ℓ is correct (see the bottom of Figure 4). A similar argument shows that the ordering of points on h_2 is correct, one needs to consider the right side of the following closed curve:

$$P_{1,\ell-1} \xrightarrow{h_{\ell-1}} P_{2,\ell-1} \xrightarrow{h_{\ell-1}} P_{\ell-1,\ell} \xrightarrow{h_{\ell-1}} B \xleftarrow{h_1} P_{1,\ell} \xleftarrow{h_1} P_{1,\ell-1}.$$

We show that $\mathcal{H}_2$ behaves similarly. Let $\mathcal{H}_2 = \{h'_1, h'_2, \ldots, h'_m\}$ and let $P'_i = g \cap h'_i$. Again, suppose that $P'_1 \xrightarrow{g} P'_2 \xrightarrow{g} \ldots \xrightarrow{g} P'_m$ (see Figure 3). By Lemma 7, h'_1 must cross $h = h_1$ at a point $Q \in A \xrightarrow{h_1} P_1$, since $h'_1 \notin \mathcal{H}_1$. Now consider the Jordan curve $e = P_1 \xrightarrow{g} P'_1 \xrightarrow{h'_1} Q \xrightarrow{h_1} P_1$. Again, it is easy to check that e separates A from P'_j for $j \geq 2$. Consider a curve h'_j ($j \geq 2$). It cannot meet h_1 before meeting g since $h'_j \notin \mathcal{H}_1$. Thus $A \xrightarrow{h'_j} P'_j$ must cross e somewhere on $P'_1 \xrightarrow{h'_1} Q$. We have reduced this problem to the previous situation with h'_1 acting as h_1, so $\mathcal{H}_2$ also forms a quasi-grid with respect to g. ☐

In the proof of the next lemma, we use the *touching graph* of a curve family of pseudo-segments. Let $\mathcal{H}$ be a family of pseudo-segments. The *touching graph* of $\mathcal{H}$ is $G_{\mathcal{H}} = (V, E)$ with $V = \mathcal{H}$ and $E = \{(a, b) : a \asymp b\}$. The statement of this lemma is almost identical to the previous one, but considers the case when the endpoints of the quasi-grid coincide. In this case, we cannot prove that the curves are the union of at most two quasi-grids, but we can still bound the number of quasi-grids by a constant.

Lemma 9. *Let $g \in \mathcal{C}(X, Y)$, and $\mathcal{H}$ is a set of g-touch-equivalent curves from $\mathcal{C}(A, A)$. If $\mathcal{H}^\star \cup \{g^\star\}$ is an intersecting family of pseudo-segments, then $\mathcal{H}$ is the disjoint union of at most 12 quasi-grids with respect to g.*

Proof. Again, we only deal with the case $g \upuparrows h$ for all $h \in \mathcal{H}$. The first claim is that any $h \in \mathcal{H}$ touches at most 2 other curves in $\mathcal{H}$. Since h is a Jordan curve, it separates the plane into two regions, one of these contains g; denote this region by R_g, and the other by R_n (see Figure 5, R_n has a line pattern). Observe that no curve in $\mathcal{H}$ touching h can enter R_n as such a curve cannot touch g. Let $P = g \cap h$. We prove that there is at most one curve in $\mathcal{H}$ that touches $A \xrightarrow{h} P$. Suppose for contradiction that curves $h_1, h_2 \in \mathcal{H}$ are both touching $A \xrightarrow{h} P$ at points T_1 and T_2 respectively, with the ordering $A \xrightarrow{h} T_1 \xrightarrow{h} T_2 \xrightarrow{h} P$. Let $Q_i = h_i \cap g$. Note that A lies on the left side of the curve $c = Q_1 \xrightarrow{g} P \xleftarrow{h} T_1 \overset{h_1}{\text{——}} Q_1$ since $g \upuparrows h$. By an earlier observation, h_2 is disjoint from the open region R_n, which is the right side of h — so h_2 touches the left side of h in T_2, i.e., $h_2 \upuparrows h$ or $h_2 \downarrow\uparrow h$. It follows that both $A \xrightarrow{h_2} T_2$ and $T_2 \xrightarrow{h_2} A$ must cross c, but this crossing can only happen along $T_1 \overset{h_1}{\text{——}} Q_1$ because the other boundary curves g and h are touched by h_2. This means that $h_2^\star$ and $h^\star$ cross at least twice, which contradicts the basic properties of an intersecting family of pseudo-segments. A similar argument shows that there is at most one curve in $\mathcal{H}$ that touches $P \xrightarrow{h} A$.

Consider the touching graph $G_{\mathcal{H}}$. Our first observation implies that the maximal degree in $G_{\mathcal{H}}$ is 2, thus by Brooks' theorem [15], $G_{\mathcal{H}}$ is 3-colorable. It is sufficient to prove that each color class is the disjoint union of at most 4 quasi-grids with respect to g.

Let $\mathcal{H}_0 \subseteq \mathcal{H}$ be a color class; consequently, it cannot contain a touching pair of curves, i.e., the curves in $\mathcal{H}$ are pairwise intersecting. Let $k = |\mathcal{H}_0|$. In this paragraph, an ending of a directed curve refers to one of the endings of its undirected version. Consider the cyclic order of the curve endings of $\mathcal{H}_0$ around A: $x_1, x_2, \ldots, x_{2k}$. Each curve appears exactly twice in this sequence. Each pair of curves in $\mathcal{H}_0$ crosses, hence x_i and x_{k+i} belong to the same curve for

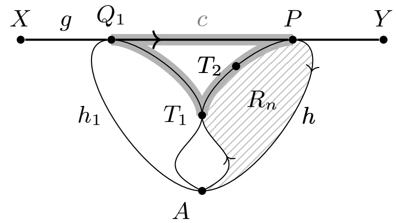

Figure 5: Curve h touches at most two other curves in $\mathcal{H}$.

each $i \in \{1, \dots, k\}$. Therefore we may dilate A to two points A_1 and A_2 such that endings $x_1, \dots, x_k$ are at A_1 and endings $x_{k+1}, \dots, x_{2k}$ are at A_2. Now $\mathcal{H}_0$ can be considered as a family of $A_1 \text{——} A_2$ curves, which is the union of $\mathcal{H}_{12}$ containing $A_1 \longrightarrow A_2$ curves and $\mathcal{H}_{21}$ containing $A_2 \longrightarrow A_1$ curves. According to Lemma 8, both $\mathcal{H}_{12}$ and $\mathcal{H}_{21}$ are the union of at most two quasi-grids with respect to g. □

The above Lemmas also imply the following one:

Lemma 10. *Let A, B, C, D be not necessarily distinct points on the plane. Let $g \in \mathcal{C}(A, B)$, and let $\mathcal{C}_0 \subset \mathcal{C}(C, D)$ be a finite curve family such that $\{g\} \cup \mathcal{C}_0$ is an intersecting family of pseudo-segments with all $h \in \mathcal{C}_0$ touching g. Then $\mathcal{C}_0$ is the disjoint union of at most 48 quasi-grids with respect to g.*

Proof. $\mathcal{C}_0$ is the disjoint union of at most four g-touching equivalence classes ($g \upuparrows h$, $g \downarrow\uparrow h$, $h \upuparrows g$ and $h \uparrow\downarrow g$). By lemmas 8 and 9, each such class can be decomposed into at most 12 quasi-grids. □

4. Touching quasi-grids with external curves

Lemma 11. *Let g be any curve in $\mathcal{C}(A, B)$ and let $\mathcal{H} = \{h_1, h_2, h_3\} \subseteq \mathcal{C}(C, D)$ be a quasi-grid with respect to g (possibly a part of a larger quasi-grid). Suppose that for a curve $g' \in \mathcal{C}(A, B)$, the set of five curves $\{g, g', h_1, h_2, h_3\}$ is an intersecting family of pseudo-segments and g' touches the middle curve $h_2 \in \mathcal{H}$. Then g' must also touch at least one more among $\{g, h_1, h_3\}$.*

Proof. Suppose that $g \upuparrows h_i$ ($i = 1, 2, 3$), the other cases are similar. The definition of quasi-grids enumerates all intersections between the four curves h_1, h_2, h_3 and g. It follows that the borders of the faces in the right side of $C \xrightarrow{h_1} P_1 \xrightarrow{g} P_3 \xrightarrow{h_3} D \xleftarrow{h_1} P_{1,3} \xleftarrow{h_3} C$ are determined (see the faces marked with encircled numbers in Figure 6). Notice that some (or all) of A, B, C and D may coincide, so the other faces are unknown. Let ① be the right side of $C \xrightarrow{h_1} P_{1,2} \xleftarrow{h_2} C$. In the same manner, we assign numbers ② – ⑤ to some other faces as well, see Figure 6.

Suppose for contradiction that g' crosses h_1, h_3 and g. We need the following claim to proceed with our proof.

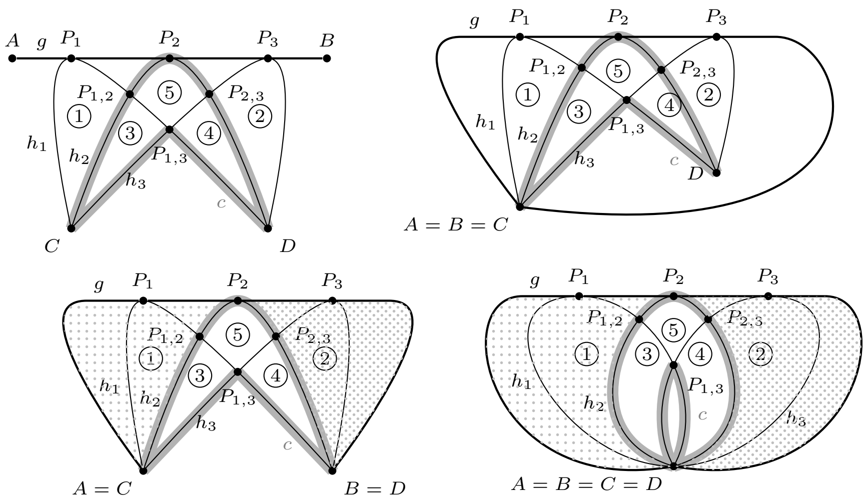

Figure 6: The figures for various possible equalities among A, B, C and D.

Claim 12. *The curve g' cannot enter region ⑤.*

Proof. If g' passes through ⑤, then – since it touches h_2 –, it must cross both $P_{1,2}\xrightarrow{h_1}P_{1,3}$ and $P_{1,3}\xrightarrow{h_3}P_{2,3}$. Since there can be no more intersections with h_1 or h_3, the curve g' cannot pass through the closed curve $c = C\xrightarrow{h_2}D\xleftarrow{h_1}P_{1,3}\xleftarrow{h_3}C$, therefore it cannot meet $g^\star$ (see Figure 6). $\square$

Since g' must cross h_1, it has to enter either region ① or ④ (by Claim 12 it cannot cross $P_{1,2}\xrightarrow{h_1}P_{1,3}$). If it enters ①, then — since it has crossed h_1, one of its endpoints A or B has to be on the border of ①, thus either $A = C$ or $B = C$. If g' enters ④, then by Claim 12, one of the endpoints is D, so $A = D$ or $B = D$. The curve g' also needs to cross h_3, so it enters either ② or ③, and as before, it follows that $A = D$ or $B = D$ in case of entering ② and $A = C$ or $B = C$ otherwise.

If g' enters ① and ③, then $A = C = B$. Since ① and ③ are on the same side of the closed curve $g \in \mathcal{C}(A, A)$, the curve $(g')^\star \in \mathcal{C}^\star(A, A)$ crosses $g^\star$ at an even number of points, we arrived at a contradiction. The case when g' enters ② and ④ is identical if we swap the role of C and D.

If g' enters ① and ②, then the endpoints of $g' \in \mathcal{C}(A, B)$ are C and D. If $A = D$ and $B = C$, then the closed curves $B\xrightarrow{h_1}P_1\xrightarrow{g}B$ and $A\xrightarrow{g}P_1\xrightarrow{h_1}A$ must cross each other at an even number of points, so there is an intersection point distinct from P_1. Note that h_1 and g are members of an intersecting family of pseudo-segments (since $\mathcal{H}$ is a quasi-grid with respect to g), so the intersection must be at their endpoints, thus $A = B$. Consequently, if g' enters ① and ②, then either $A = B = C = D$ or $A = C$ and $B = D$ are two distinct points

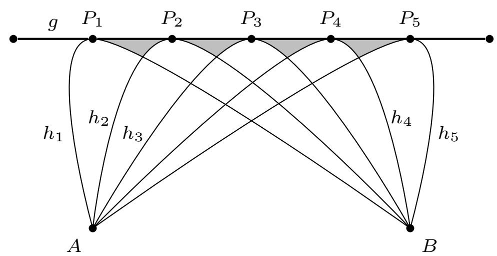

Figure 7: A 5-element quasi-grid $\mathcal{H}'$

(see the bottom of Figure 6). Let R_1 be the region to the left of $A \xrightarrow{h_2} P_2 \xleftarrow{g} A$ (sparsely dotted) and R_2 be the left side of $B \xleftarrow{g} P_2 \xrightarrow{h_2} B$ (densely dotted). Notice that g' starts in R_1 and ends in R_2, two regions that are guaranteed to be disjoint apart from P_2, A and B. Since it cannot cross h_2, it crosses both $g_1^\star$ and $g_2^\star$, where $g_1 = A \xrightarrow{g} P_2$ and $g_2 = P_2 \xrightarrow{g} B$. This is a contradiction since g' has to cross g exactly once.

If g' enters ③ and ④, then $A = C$ and $B = D$ by the same argument as in the previous case. Let A' and B' be points on g' close to its starting and endpoint A and B, so that there are no touchings or crossings on g' between A and A' and between B' and B. Note that g' cannot cross h_2 because they need to touch. Thus the boundary of R_1 can only be crossed on g_1, and both A' and B' lie outside R_1 (they are in ③ and ④ respectively), so the number of intersections between $g_1^\star$ and $(g')^\star$ is even. The same argument holds for R_2 and g_2, so the number of intersections between $g^\star$ and $(g')^\star$ is even – a contradiction. □

The next lemma demonstrates our intuitive claim that touching the members of a large quasi-grid by two curves is not possible inside an intersecting family of pseudo-segments.

Lemma 13. *Let $\mathcal{H}$ be a set of at least 5 curves from $\mathcal{C}(A, B)$, where A and B may coincide. Let g_1, g_2 be two curves such that $\mathcal{H} \cup \{g_1, g_2\}$ form a family of pseudo-segments. Then $\mathcal{H}$ cannot form a quasi-grid with respect to both g_1 and g_2.*

Proof. Suppose for contradiction that $\mathcal{H}$ is a quasi-grid with respect to g_1 and g_2 simultaneously. Let $\mathcal{H}' = \{h_1, h_2, \ldots, h_5\}$ be a 5-element subset of $\mathcal{H}$ that touch g_1 in this order at $P_1, \ldots, P_5$, see Figure 7.

The curve g_2 cannot have any points in a region which is enclosed by only curves from $\mathcal{H}'$: it cannot leave the region since it cannot cross any of $\mathcal{H}'$, and every region is bounded by at most four of the $\mathcal{H}'$ curves, so at least one curve would remain untouchable for g_2.

Consequently, g_2 has to touch h_2, h_3 and h_4 in the regions enclosed by g_1, h_i and h_{i+1} ($i = 1, 2, 3, 4$) (see the shaded regions in Figure 7). Since g_2 can meet

g_1 at most once, it can visit only one of these regions, so at least one of h_2, h_3 and h_4 will remain untouchable - we arrived at a contradiction.

□

Acknowledgements

We thank János Pach and Géza Tóth for suggesting the original problem, for the encouragement and for the fruitful discussions. We thank an anonymus referee for several remarks that improved the presentation of the paper. This research was supported by grant no. K 109240 from the National Development Agency of Hungary, based on a source from the Research and Technology Innovation Fund.

References

[1] M. Ajtai, V. Chvátal, M. M. Newborn, E. Szemerédi, Crossing-free subgraphs, North-Holland Mathematics Studies 60 (1982) 9–12. `doi:10.1016/S0304-0208(08)73484-4`.

[2] F. T. Leighton, Complexity Issues in VLSI: Optimal Layouts for the Shuffle-exchange Graph and Other Networks, MIT Press, Cambridge, MA, USA, 1983.

[3] E. Szemerédi, W. T. Trotter Jr, Extremal problems in discrete geometry, Combinatorica 3 (3-4) (1983) 381–392. `doi:10.1007/BF02579194`.

[4] K. Kedem, R. Livne, J. Pach, M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, Discrete Comput. Geom. 1 (1) (1986) 59–71. `doi:10.1007/BF02187683`.

[5] H. Tamaki, T. Tokuyama, How to cut pseudoparabolas into segments, Discrete Comput. Geom. 19 (2) (1998) 265–290. `doi:10.1007/PL00009345`.

[6] B. Aronov, M. Sharir, Cutting circles into pseudo-segments and improved bounds for incidences, Discrete Comput. Geom. 28 (4) (2002) 475–490. `doi:10.1007/s00454-001-0084-1`.

[7] P. K. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir, S. Smorodinsky, Lenses in arrangements of pseudo-circles and their applications, J. ACM 51 (2) (2004) 139–186. `doi:10.1145/972639.972641`.

[8] A. Marcus, G. Tardos, Intersection reverse sequences and geometric applications, J. Combin. Theory Ser. A 113 (4) (2006) 675–691. `doi:10.1016/j.jcta.2005.07.002`.

[9] J. Fox, J. Pach, C. D. Tóth, Intersection patterns of curves, J. London Math. Soc. 83 (2011) 389–406. `doi:10.1112/jlms/jdq087`.

[10] R. B. Richter, C. Thomassen, Intersections of curve systems and the crossing number of $C_5 \times C_5$, Discrete Comput. Geom. 13 (1) (1995) 149–159. doi:10.1007/BF02574034.

[11] J. Pach, N. Rubin, G. Tardos, Beyond the Richter-Thomassen conjecture, in: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, 2016, pp. 957–968. doi:10.1137/1.9781611974331.ch68.

[12] J. Fox, F. Frati, J. Pach, R. Pinchasi, Crossings between curves with many tangencies, An Irregular Mind, Bolyai Soc. Math. Stud. 21 (2010) 251–260. doi:10.1007/978-3-642-14444-8_6.

[13] J. Pach, Private communication (2015).

[14] J. Pach, M. Sharir, On vertical visibility in arrangements of segments and the queue size in the bentley-ottmann line sweeping algorithm, SIAM Journal on Computing 20 (3) (1991) 460–470. doi:10.1137/0220029.

[15] R. L. Brooks, On colouring the nodes of a network, Proc. Cambridge Philosophical Society, Math. Phys. Sci. 37 (1941) 194–197. doi:10.1017/S030500410002168X.