Extremal Configurations and Levels in Pseudoline Arrangements *

Micha Sharir[†] Shakhar Smorodinsky[‡]

Abstract

This paper studies a variety of problems involving certain types of extreme configurations in arrangements of (x-monotone) pseudo-lines. For example, we obtain a very simple proof of the bound $O(nk^{1/3})$ on the maximum complexity of the k-th level in an arrangement of n pseudo-lines, which becomes even simpler in the case of lines. We thus simplify considerably previous proofs by Dey and by Tamaki and Tokuyama.

We also consider *diamonds* and *anti-diamonds* in (simple) pseudo-line arrangements, where a diamond (resp., an anti-diamond) is a pair u, v of vertices, so that u lies in the double wedge of v and vice versa (resp., neither u nor v lies in the other double wedge). We show that the maximum size of a diamond-free set of vertices in an arrangement of n pseudo-lines is $3n - 6$, by showing that the induced graph (where each vertex of the arrangement is regarded as an edge connecting the two incident curves) is planar, simplifying considerably a previous proof of the same fact by Tamaki and Tokuyama. Similarly, we show that the maximum size of an anti-diamond-free set of vertices in an arrangement of n pseudo-lines is $2n-2$. We also obtain several additional results, which are listed in the introduction.

In some of our results, we use a recent duality transform between points and pseudo-lines due to Agarwal and Sharir, which extends an earlier transform by Goodman (that applied only in the projective plane). We show that this transform maps a set of vertices in a pseudo-line arrangement to a topological graph whose edges are drawn as x-monotone arcs that connect pairs of the dual points, and form a set of extendible pseudo-segments (they are pieces of curves that form a pseudo-line arrangement in the dual plane). This allows us (a) to 'import' known results on this kind of topological graphs to the context of pseudo-lines; (b) to extend techniques that have been originally applied only for geometric graphs (whose edges are drawn as straight segments), thereby obtaining new results for pseudo-line arrangements, or for the above-class of x-monotone topological graphs; and (c) to derive new techniques, facilitated by the passage to the dual setting, that apply in the more general pseudo-line context, and extend and simplify considerably the earlier proofs. This paper contains examples of all three kinds of results.

1 Introduction

Let Γ be a collection of n *pseudolines* in the plane, which we define to be graphs of continuous totally-defined functions, each pair of which intersect in exactly one point, and the curves cross each other at that point. In what follows we assume gen-

*Work on this paper has been supported by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing). Work by Micha Sharir has also been supported by NSF Grants CCR-97-32101 and CCR-00-98246, by a grant from the U.S.-Israeli Binational Science Foundation, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

[†]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; sharir@math.tau.ac.il.

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; smoro@post.tau.ac.il. The research by this author was done while the author was a Ph.D. student under the supervision of Micha Sharir.

eral position of the pseudolines, meaning that no three pseudolines pass through a common point, and that the x-coordinates of any two intersection points of the pseudolines are distinct. Let E be a subset of the vertices of the arrangement $A(\Gamma)$. E induces a graph $G = (\Gamma, E)$ on Γ (in what follows, we refer to such a graph as a *pseudoline graph*). For each pair (γ, γ') of distinct pseudolines in Γ, we denote by $W(\gamma, \gamma')$ the *double wedge* formed between γ and γ', that is, the (open) region consisting of all points that lie above one of these pseudolines and below the other. We also denote by $W^c(\gamma, \gamma')$ the complementary (open) double wedge, consisting of all points that lie either above both curves or below both curves.

Definition 1.1 *We say that two edges* (γ, γ') *and* (δ, δ') *of* G *form a* diamond *if the point* $\gamma \cap \gamma'$ *is contained in the double wedge* $W(\delta, \delta')$*, and the point* $\delta \cap \delta'$ *is contained in the double wedge* $W(\gamma, \gamma')$*. See e.g. Figure 3.*

Definition 1.2 *We say that two edges* (γ, γ') *and* (δ, δ') *of* G *form an* anti-diamond *if the point* $\gamma \cap \gamma'$ *is not contained in the double wedge* $W(\delta, \delta')$*, and the point* $\delta \cap \delta'$ *is not contained in the double wedge* $W(\gamma, \gamma')$*; that is,* $\gamma \cap \gamma'$ *lies in* $W^c(\delta, \delta')$ *and* $\delta \cap \delta'$ *lies in* $W^c(\gamma, \gamma')$*. See e.g. Figure 6.*

Definition 1.3 *(a) A collection* S *of* x*-monotone bounded Jordan arcs is called a collection of* pseudosegments *if each pair of arcs of* S *intersect in at most one point, where they cross each other.*
(b) S *is called a collection of* extendible pseudosegments *if there exists a set* Γ *of pseudolines, with* $|\Gamma| = |S|$*, such that each* $s \in S$ *is contained in a unique pseudoline of* Γ*.*

See [8] for more details concerning extendible pseudosegments. Note that not every collection of pseudosegments is extendible, as shown by the simple example depicted in Figure 1.

Definition 1.4 *(a) A* drawing *of a graph* $G = (\Gamma, E)$ *in the plane is a mapping that maps each vertex* $v \in \Gamma$ *to a point in the plane, and each edge* $e = uv$ *of* E *to a Jordan arc connecting the images of* u *and* v*, such that no three arcs are concurrent at their relative interiors, and the relative interior of no arc is incident to a vertex.*

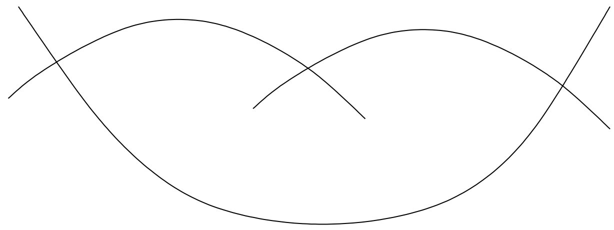

Figure 1: Three pseudosegments that are not extendible.

(b) If the images of the edges of E *form a family of extendible pseudo-segments then we refer to the drawing of* G *as an (*x*-monotone)* generalized geometric graph.

(The term *geometric graphs* is usually reserved to drawings of graphs where the edges are drawn as straight segments.)

In this paper we prove the following results.

Duality between pseudoline graphs and generalized geometric graphs. The first main result of this paper establishes an equivalence between pseudoline graphs and generalized geometric graphs.

We first derive the following weaker result, which has an easy and self-contained proof.

Theorem 1.5 *Let* Γ *and* G *be as above. Then there is a drawing of* G *in the plane such that two edges* e *and* e' *of* G *form a diamond if and only if their corresponding drawings cross each other an odd number of times.*

After the original preparation of this paper, Agarwal and Sharir [5] established a duality transformation in arrangements of pseudolines, which has several useful properties and other applications. Using their technique, we derive the following stronger result:

Theorem 1.6 *(a) Let* Γ *and* G *be as above. Then there is a drawing of* G *in the plane, with the edges constituting a family of extendible pseudosegments, such that, for any two edges* e*,* e' *of* G*,* e *and* e' *form a diamond if and only if their corresponding drawings cross each other.*
(b) Conversely, for any graph $G = (V, E)$ *drawn in the plane with its edges constituting a family of extendible pseudosegments, there exists a family* Γ *of pseudolines and a 1-1 mapping* φ *from* V *onto*

Γ, *so that each edge* $uv \in E$ *is mapped to the vertex* $\varphi(u) \cap \varphi(v)$ *of* $\mathcal{A}(\Gamma)$, *such that two edges in* E *cross each other if and only if their images are two vertices of* $\mathcal{A}(\Gamma)$ *that form a diamond.*

Applications. As an immediate corollary of Theorem 1.6 (which can also be derived from Theorem 1.5, using the fact [25] that any graph drawn in the plane such that every pair of edges on four distinct vertices cross an even number of times, is planar), we obtain

Theorem 1.7 *Let* Γ *and* G *be as above. If* G *is diamond-free then* G *is planar and thus* $|E| \leq 3n-6$.

Theorem 1.7 has been proven by Tamaki and Tokuyama [23], using a more involved argument. This was the underlying theorem that enabled them to extend Dey's improved bound of $O(n^{4/3})$ on the complexity of a single level in an arrangement of lines [11], to arrangements of pseudolines. Note that the planarity of G is obvious for the case of lines: If we dualize the given lines into points, using the duality $y = ax + b \mapsto (a, b)$ and $(c, d) \mapsto y = -cx + d$, presented in [13], and map each edge (γ, γ') of G to the straight segment connecting the points dual to γ and γ', we obtain a crossing-free drawing of G. Hence, Theorem 1.7 is a natural (though harder to derive) extension of this property to the case of pseudolines.

We note also that the converse statement of Theorem 1.7 is trivial: Every planar graph can be realized as a diamond-free pseudoline graph (in fact, in an arrangement of lines): We draw the graph as a straight-edge graph (which is always possible [14]), and apply the inverse duality to the one just mentioned.

In more generality, we can take any theorem that involves generalized geometric graphs (whose edges are extendible pseudosegments), and that studies the crossing pattern of these edges, and 'transport' it into the domain of pseudoline graphs. As an example of this, we have:

Theorem 1.8 *Let* Γ *and* G *be as above. (i) If* G *contains no three edges which form pairwise diamonds then* G *is quasi-planar (in the terminology of [2]; see below), and thus its size is* $O(n)$. *(ii) If* G *contains no* k *edges which form pairwise diamonds (for any fixed* $k \geq 4$*) then the size of* G *is* $O(n \log n)$ *(with the constant of proportionality depending on* k*).*

In its appropriate reformulation in the context of generalized geometric graphs, Theorem 1.8(i) corresponds to a result of Agarwal et al. [2] on quasi-planar graphs. A *quasi-planar* (respectively, k-*quasi-planar*) graph is a graph that can be drawn in the plane such that no three (respectively, k) of its edges are pairwise crossing. It was shown in [2] that the size of a quasi-planar graph is $O(n)$. This result was extended by Valtr [26] to the case $k \geq 4$ and our Theorem 1.8(ii) is a similar interpretation of Valtr's bound in the context of pseudoline graphs. Our reformulations are valid, for both parts of the theorem, since both the results of [2, 27] hold for graphs whose edges are extendible pseudosegments.

Definition 1.9 *A* thrackle *is a drawing of a graph in the plane so that every pair of edges either have a common endpoint and are otherwise disjoint, or else they intersect in exactly one point where they cross each other.*

The notion of a thrackle is due to Conway, who conjectured that the number of edges in a thrackle is at most the number of vertices. The study of thrackles has drawn much attention. Two recent papers [18] and [7] obtain linear bounds for the size of a general thrackle, but with constants of proportionality that are greater than 1. The conjecture is known to hold for straight-edge thrackles [20], and, in Section 6, we extend the result, and the proof, to the case of graphs whose edges are extendible pseudosegments. That is, we show:

Theorem 1.10 *Let* Γ *and* G *be as above. If every pair of edges connecting four distinct vertices (that is, curves of* Γ*) in* G *form a diamond, then the size of* G *is at most* n.

Our proof extends ideas used by Perles in the proof for the straight edge case.

Pseudoline graphs without anti-diamonds. We now turn to study pseudoline graphs that do not have any anti-diamond. We show:

Theorem 1.11 *Let* Γ *and* G *be as above. If* G *is anti-diamond-free then* $|E| \leq 2n-2$.

Theorem 1.11 is an extension, to the case of pseudolines, of a (dual version of a) theorem of Katchalski and Last [15], refined by Valtr [27], both solving a problem posed by Kupitz. The theorem states that a straight-edge graph on n points in the plane, which does not have any pair of parallel edges, has at most $2n-2$ edges. A pair of segments e, e' is said to be *parallel* if the line containing e does not cross e' and the line containing e' does not cross e. (For straight edges, this is equivalent to the condition that e and e' are in convex position.) The dual version of a pair of parallel edges is a pair of vertices in a line arrangement that form an anti-diamond. Hence, Theorem 1.11 is indeed an extension of the result of [15, 27] to the case of pseudolines. The proof, for the case of straight-edge graphs, has been recently simplified by Valtr [28]. Our proof, obained independently, can be viewed as an extension of this new proof to the case of pseudolines.

Note that Theorem 1.11 is not directly obtainable from [15, 27, 28], (a) because Theorem 1.6 does not cater to anti-diamonds, and (b) because the analysis of [15, 27, 28] only applies to straight-edge graphs.

The complexity of the k-level in an arrangement of pseudolines. We provide a simpler proof of the following major result in combinatorial and computational geometry:

Theorem 1.12 *The maximum complexity of the k-level in an arrangement of n pseudolines is $O(nk^{1/3})$.*

The k-level in the arrangement of a set Γ of n pseudolines is the (closure of) the set of all points that lie on curves of Γ and have exactly k other curves passing below them. This is a central structure in arrangements, with a long and rich history, and with many applications, both in discrete and in computational geometry; See e.g., [19]. In a recent breakthrough, Dey [11] has shown that the complexity (number of vertices) of the k-level in an arrangement of n lines is $O(nk^{1/3})$ (the best known lower bound is only near-linear [24]). This bound was extended to the case of pseudolines by Tamaki and Tokuyama [23], using a very complicated proof. We present a much simpler proof (than both proofs in [11] and [23]) for the general case of pseudolines.

Incidences and many faces in pseudoline arrangements. Finally, as an application of Theorem 1.7, we provide yet another simple proof of the following well-known result in a much-studied research area:

Theorem 1.13 *(a) The maximum number of incidences between m distinct points and n distinct pseudolines is $\Theta(m^{2/3}n^{2/3} + m + n)$.*

(b) The maximum number of edges bounding m distinct faces in an arrangement of n pseudolines is $\Theta(m^{2/3}n^{2/3} + n)$.

The proof is in some sense 'dual' to the proofs based on Székely's technique [12, 22].

The proof of Theorem 1.13(b) can be extended to yield the following result, recently obtained in [3], where it has been proved using the dual approach, based on Székely's technique.

Theorem 1.14 *The maximum number of edges bounding m distinct faces in an arrangement of n extendible pseudo-segments is $\Theta((m+n)^{2/3}n^{2/3})$.*

2 Drawing Pseudoline Graphs

In this section we prove Theorems 1.5 and 1.6. Both proofs use the same drawing rule for realizing pseudoline graphs as geometric graphs. The difference is that the stronger properties of Theorem 1.6 follow from the more sophisticated machinery of point-pseudoline duality, developed in [5]. On the other hand, the proof of Theorem 1.5 is simple and self-contained.

Proof of Theorem 1.5: Let ℓ be a vertical line such that all vertices of the arrangement $A(\Gamma)$ lie to the right of ℓ. Enumerate the pseudolines of Γ as $\gamma_1, \ldots, \gamma_n$, ordered in increasing y-coordinates of the intersection points $p_i = \ell \cap \gamma_i$. We construct a drawing of G in the plane, using the set $P = \{p_1, \ldots, p_n\}$ as the set of vertices.

For each edge $(\gamma_i, \gamma_j) \in E$, we connect the points p_i and p_j by a y-monotone curve $e_{i,j}$ according to the following rules. Assume, without loss of generality, that $i > j$. If $i = j+1$ (so that p_i and p_j are consecutive intersection points along ℓ) then $e_{i,j}$ is just the straight segment p_ip_j (contained in ℓ). Otherwise, $e_{i.j}$ is drawn very close to ℓ, and generally proceeds upwards (from p_j to p_i) parallel to ℓ either slightly to its left or slightly to its right. In the vicinity of an intermediate point p_k, the edge either continues parallel to ℓ, or converges to p_k (if $k = i$), or switches to the other side of ℓ, crossing it before p_k. The decision on which side of p_k the edge should pass is made according to the following

> **Drawing rule:** If the pseudoline γ_k passes above the apex of $W(\gamma_i, \gamma_j)$ then $e_{i,j}$ passes to the left of p_k, otherwise $e_{i,j}$ passes to the right of p_k.

This drawing rule is a variant of a rule recently proposed in [4] for drawing, and proving the planarity, of another kind of graphs related to arrangements of pseudocircles or pseudo-parabolas. Note that this rule does not uniquely define the drawing.

We need the following technical lemma:

Lemma 2.1 *Let $x_1 < x_2 < x_3 < x_4$ be four real numbers. (i) Let $e_{1,4}$ and $e_{2,3}$ be two x-monotone Jordan arcs with endpoints at $(x_1, 0), (x_4, 0)$ and $(x_2, 0), (x_3, 0)$, respectively, so that $e_{1,4}$ does not pass through $(x_2, 0)$ or through $(x_3, 0)$. Then $e_{1,4}$ and $e_{2,3}$ cross an odd number of times if and only if $e_{1,4}$ passes around the points $(x_2, 0)$ and $(x_3, 0)$ on different sides. See Figure 2(a).*

(ii) Let $e_{1,3}$ and $e_{2,4}$ be two x-monotone Jordan arcs with endpoints at $(x_1, 0), (x_3, 0)$ and $(x_2, 0), (x_4, 0)$, respectively, so that $e_{1,3}$ does not pass through $(x_2, 0)$ and $e_{2,4}$ does not pass through $(x_3, 0)$. Then $e_{1,3}$ and $e_{2,4}$ cross an odd number of times if and only if $e_{1,3}$ passes below $(x_2, 0)$ and $e_{2,4}$ passes below $(x_3, 0)$, or $e_{1,3}$ passes above $(x_2, 0)$ and $e_{2,4}$ passes above $(x_3, 0)$. See Figure 2(b).

Proof: In case (i), let f_1 and f_2 be the two real (partially defined) continuous functions whose graphs are $e_{1,4}$ and $e_{2,3}$, respectively. Similarly, for case (ii), let f_1 and f_2 be the functions whose graphs are $e_{1,3}$ and $e_{2,4}$, respectively.

Consider the function $g = f_1 - f_2$ over the interval $[x_2, x_3]$. By the mean-value theorem, $g(x_2)$ and $g(x_3)$ have different signs if and only if g vanishes an odd number of times over this interval. This completes the proof of the Lemma. $\square$

Let $e_1 = e_{x,y}$, $e_2 = e_{z,w}$ be the drawings of two distinct edges in G that do not share a vertex. We consider two possible cases:

Case (i): The intervals $p_x p_y$ and $p_z p_w$ (on the line ℓ) are nested. That is, their endpoints are ordered, say, as p_z, p_x, p_y, p_w in y-increasing order along the line ℓ. By Lemma 2.1, e_1 and e_2 cross an odd number of times if and only if e_2 passes around the points p_x and p_y on different sides. On the other hand, it is easily checked that the drawing rule implies that e_1 and e_2 form a diamond in G if and only if e_2 passes around the points p_x and p_y on different sides. Hence, in this case we have that e_1 and e_2 form a diamond if and only if they cross an odd number of times. See Figure 3 for an illustration.

Case (ii): The intervals $p_x p_y$ and $p_z p_w$ 'interleave', so that the y-order of the endpoints of e_1 and e_2 is, say, p_x, p_z, p_y, p_w, or a symmetrically similar order. By Lemma 2.1, e_1 and e_2 cross an odd number of times if and only if e_1 passes around the point p_z on the same side that e_2 passes around the point p_y. On the other hand, the drawing rule for e_1 and e_2 easily implies that e_1 and e_2 form a diamond if and only if e_1 passes around the point p_z on the same side that e_2 passes around the point p_y. See Figure 4 for an illustration.

It is also easily checked that, in the case where the intervals $p_x p_y$ and $p_z p_w$ are disjoint, the edges e_1 and e_2 do not form a diamond, nor can their drawings intersect each other. This completes the proof of the theorem. $\square$

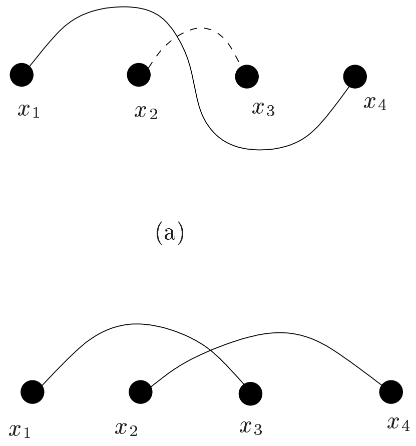

(a)

(b)

Figure 2: Two instances where a pair of drawn edges have an odd number of crossings. (a) The nested case. (b) The interleaving case.

Proof of Theorem 1.6: The drawing rule used in the proof of Theorem 1.5 is in fact a special case of the duality transform between points and (x-monotone) pseudolines, as obtained recently by Agarwal and Sharir [5]. Specifically, we apply this result to Γ and to the set G of the given vertices of $\mathcal{A}(\Gamma)$. The duality of [5] maps the points of G to a set G^* of x-monotone pseudolines, and maps the pseudolines of Γ to a set Γ^* of points, so that a point $v \in G$ lies on (resp., above, below) a curve $\gamma \in \Gamma$ if and only if the dual pseudoline v^* passes through (resp., above, below) the dual point γ^*. Finally, in the transformation of [5], the points of Γ^* are arranged along the x-axis in the same order as that of the intercepts of these curves with the vertical line ℓ defined above.

We apply this transformation to Γ and G. In addition, for each vertex $v \in G$, incident to two pseudolines $\gamma_1, \gamma_2 \in \Gamma$, we trim the dual pseudoline v^* to its portion between the points γ_1^*, γ_2^*.

This yields a plane drawing of the graph G, whose edges form a collection of extendible pseudo-segments. The drawing has the following main property:

Lemma 2.2 *Let $v = \gamma_1 \cap \gamma_2$ and $w = \gamma_3 \cap \gamma_4$ be two vertices in G, defined by four distinct curves. Then v and w form a diamond if and only if the corresponding edges of the drawing cross each other.*

Proof: The proof is an easy consequence of the proof of Theorem 1.5 given above. In fact, it suffices to show that the duality transformation of [5] obeys the drawing rule used in the above proof, with an appropriate rotation of the plane by 90 degrees. So let $\gamma_i, \gamma_j, \gamma_k \in \Gamma$ such that γ_k passes above (resp., below) $\gamma_i \cap \gamma_j$, and such that γ_k meets the vertical line ℓ at a point between $\gamma_i \cap \ell$ and $\gamma_j \cap \ell$. Our drawing rule then requires that the edge $p_i p_j$ pass to the left (resp., to the right) of p_k. On the other hand, the duality transform, preserving the above/below relationship, makes the edge $\gamma_i^* \gamma_j^*$ pass below (resp., above) γ_k^*. Hence the two rules coincide, after an appropriate rotation of the plane, and the lemma is now an easy consequence of the preceding analysis. $\square$

Lemma 2.2 thus implies Theorem 1.6(a). To prove the converse part (b), let $G = (V, E)$ be a graph drawn in the plane so that its edges form a collection of extendible pseudo-segments, and let Λ denote the family of pseudolines containing the edges of E. Apply the point-pseudoline duality transform of [5] to V and Λ. We obtain a family V^* of pseudolines and a set Λ^* of points, so that the incidence and the above/below relations between V and Λ are both preserved. It is now routine to verify, as in the case of point-line duality, that two edges u_1v_1 and u_2v_2 of E cross each other if and only if the corresponding vertices $u_1^* \cap v_1^*$, $u_2^* \cap v_2^*$ of $\mathcal{A}(V^*)$ form a diamond. This completes the proof of Theorem 1.6. $\square$

The immediate implications of these results, namely Theorems 1.7 and 1.8, follow as well, as discussed in the introduction.

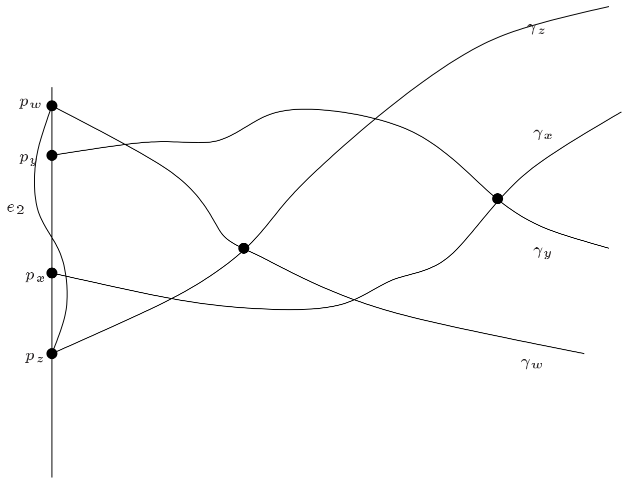

Figure 3: A diamond, and the resulting crossing in the case that the segments p_xp_y and p_zp_w are nested.

3 The Complexity of a k-Level in Pseudoline Arrangements

In this section we provide a simpler proof of the well-known $O(nk^{1/3})$ upper bound on the complexity of the k-level in an arrangement of n pseudolines (see [11, 23]).

Let Γ be the given collection of n pseudolines, and let E be the set of vertices of the k-level, where k is a fixed constant ($0 \le k \le n-2$). Theorem 1.6 and a standard probabilistic argument allow us to derive the following extension of the Crossing Lemma of [6, 16]; We omit the proof here.

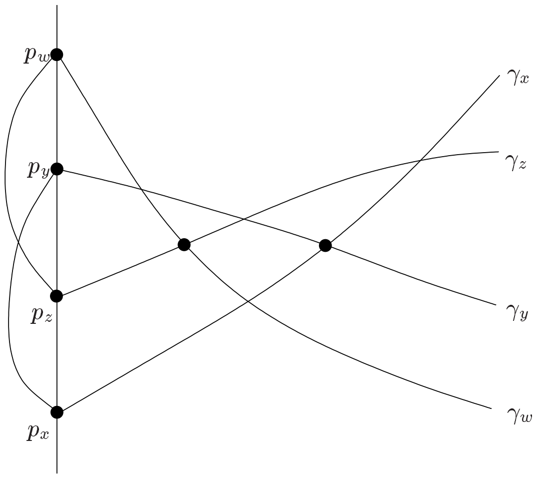

Figure 4: A diamond, and the resulting crossing in the case that the segments p_xp_y and p_zp_w are interleaved.

Lemma 3.1 (Extended Crossing Lemma) *Let $G(\Gamma, E)$ be a pseudoline graph on n pseudolines, with $|E| \geq 4n$. Then G has $\Omega(|E|^3/n^2)$ diamonds.*

Remark: In the restricted case where Γ is a collection of real lines, Lemma 3.1 is a dual version of the Crossing Lemma of [6, 16].

Dey [11] has shown that the number of diamonds in G is at most the total number of vertices in the arrangement $A(\Gamma)$ that lie at level less than k. It is well known (see e.g. [10]) that the overall complexity of the first k levels in an arrangement of n lines or pseudolines is $O(nk)$. Hence, this fact, combined with the lower bound discussed above, yield the $O(nk^{1/3})$ upper bound on the complexity of the k-level. We provide here an alternative simpler proof that the number of diamonds in G is at most $k(n-k-1)$, without using the bound on the complexity of the first k levels.

We use the fact that the vertices of the k-level can be grouped into vertex sets of k 'concave' vertex-disjoint chains $c_1, \ldots, c_k$. Each such chain c_i is an x-monotone (connected) path that is contained in the union of the pseudolines of Γ, such that all the vertices of c_i are at level k. Informally, as we traverse c_i from left to right, whenever we reach a vertex of $A(\Gamma)$, we can either continue along the pseudoline we are currently on, or make a right (i.e., downward) turn onto the other pseudoline, but we cannot make a left (upward) turn; in case the pseudolines are real lines, c_i is indeed a concave polygonal chain. The simple construction of these chains is described in [1]: Each chain starts on each of the k lowest pseudolines at $x = -\infty$, and makes (right) turns *only* at vertices of the k-level. In a symmetric manner we can group the vertices of the k-level into $n-k-1$ 'convex' vertex-disjoint chains, by starting the chains along the $n-k-1$ highest pseudolines at $x = -\infty$, and by making *left* turns only at vertices of the k-level. Let (p, p') be a diamond, where p and p' are vertices at level k. Assume without loss of generality that p lies to the left of p'. Let c_1 be the unique *concave* chain that contains p and let c_2 be the unique *convex* chain that contains p'. See Figure 5 for an illustration. For a given vertex v in $A(\Gamma)$, let $W_r(v)$ (resp. $W_l(v)$) denote the (interior of the) right (resp. left) wedge of the double-wedge formed by the two pseudolines defining v. Consider the right wedges of vertices of c_1. It is easy to see (from the 'concavity' of c_1) that those wedges are pairwise disjoint (see also [1]). A similar argument holds for the left wedges of the vertices of c_2. Since $p' \in W_r(p)$ and $p \in W_l(p')$, it follows that c_2 does not meet the lower edge of $W_r(p)$, but meets the upper edge of this wedge. This can happen for at most one vertex of c_1, because of the disjointness of the right wedges of its vertices. Hence p is uniquely determined from the pair (c_1, c_2), and, symmetrically this also holds for p'.

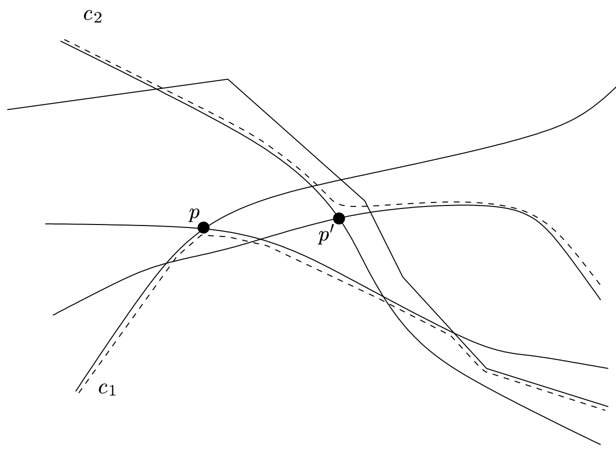

Figure 5: An arrangement of 5 pseudolines and a diamond formed by two vertices (p, p') at level 1, the corresponding concave chain c_1 containing p, and the convex chain c_2 containing p'.

Thus the number of diamonds in the k-level is at most the number of pairs (c_1, c_2) of a concave chain and a convex chain; that is, at most $k(n-k-1)$. This completes the proof of Theorem 1.12.

4 Yet Another Proof for Incidences and Many Faces in Pseudoline Arrangements

In this section we provide yet another proof of the well-known (worst-case tight) bounds given in Theorem 1.13.

We will prove only part (b) of the theorem; part (a) can then be obtained by a simple and known reduction (see, e.g., [9]); alternatively, it can be obtained by a straightforward modification of the proof of (b), given below.

Let Γ be the given collection of n pseudolines, and let $f_1, \ldots, f_m$ be the m given faces of the arrangement $\mathcal{A}(\Gamma)$. Let E denote the set of all vertices of these faces, excluding the leftmost and rightmost vertex, if any, of each face. Since every bounded face has at least one vertex that is not leftmost or rightmost, and since the number of unbounded faces is $O(n)$, it follows that the quantity that we wish to bound is $O(|E|+n)$. By Lemma 3.1, if $|E| \geq 4n$ then the graph $G(\Gamma, E)$ has $\Omega(|E|^3/n^2)$ diamonds. Let (p, p') be a diamond, where p is a vertex of some face f and p' is a vertex of another face f'. (It is easily verified that if p and p' bound the same face then they cannot form a diamond.) Then, using the Levy Enlargement Lemma [17], there exists a curve γ_0 that passes through p and p', such that $\Gamma \cup \{\gamma_0\}$ is still a family of pseudolines. In this case γ_0 must be contained in the two double wedges of p and p', and thus it avoids the interiors of f and of f'; that is, γ_0 is a 'common tangent' of f and f'. As in the case of lines, it is easy to show that a pair of faces can have at most four common tangents of this kind. Hence, the number of diamonds in G cannot exceed $2m^2$. Putting everything together, we obtain $|E| = O(m^{2/3}n^{2/3}+n)$. $\square$

Remark: This proof is, in a sense, dual to that of Székely [22] for incidences, or to its extension by Dey and Pach [12] for many faces. These former proofs interchange the roles of points and (pseudo)lines: they apply the crossing lemma to a different graph, whose vertices are the points involved in the incidences or marking points, one in each of the given faces.

The proof of Theorem 1.14 is proved in a similar manner. The main difference is that the given faces need not be x-monotone, because their boundaries may contain endpoints of the given pseudo-segments. In this case two vertices of the same face may form a diamond, and the number of diamonds formed between two distinct faces may be arbitrarily large. To overcome this issue, we partition, as in [3], any such face into x-monotone subfaces, by vertical segments erected from endpoints of the pseudo-segments. The number of new subfaces is $O(m+n)$, and any pair of them can induce only $O(1)$ diamonds, which can be argued exactly as in the case of pseudolines. The preceding arguments then yield the asserted bound.

5 Graphs in Pseudoline Arrangements without Anti-Diamonds

So far, the paper has dealt exclusively with the existence or nonexistence of diamonds in graphs in pseudoline arrangements. We now turn to graphs in pseudoline arrangements that do not contain any *anti-diamond*. Recall that the notion of an anti-diamond is an extension, to the case of pseudolines, of (the dual version of) a pair of edges in (straight-edge) geometric graphs that are in convex position (so-called 'parallel' edges). Using Theorem 1.6 (and the analysis in its proof), one obtains a transformation that maps an anti-diamond-free pseudoline graph (Γ, G) to a generalized geometric graph, whose edges form a collection of extendible pseudo-segments, with the property that, for any pair e, e' of its edges, defined by four distinct vertices, either the pseudoline containing e crosses e' or the pseudoline containing e' crosses e.

We present a much shorter and simpler proof of Theorem 1.11 than those of [15, 27], that applies directly in the original pseudoline arrangement, and is similar in spirit to the recent simplified proof of Valtr [28] for the case of straight-edge geometric graphs.

Proof of Theorem 1.11: We construct two sequences A and B whose elements belong to Γ, as follows. We sort the intersection points of the pseudolines of Γ that correspond to the edges of G in increasing x-order, and denote the sorted sequence

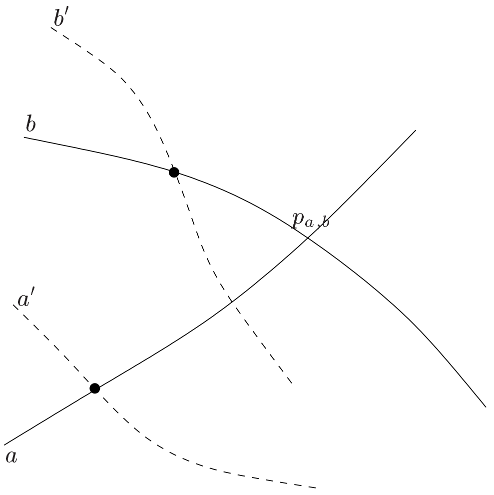

Figure 6: A sub-sequence $\cdots a \cdots b \cdots$ of A to the left of $p_{a,b}$ and the resulting anti-diamond.

by $P = \langle p_1, \ldots, p_m \rangle$. For each element p_i of P, let γ_i and γ_i' be the two pseudolines forming (meeting at) p_i, so that γ_i lies below γ_i' to the left of p_i (and lies above γ_i' to the right). Then the i-th element of A is γ_i and the i-th element of B is γ_i'.

Lemma 5.1 *The concatenated cyclic sequence $C = A\|B$ does not contain a subcycle of alternating symbols of the form $a \cdots b \cdots a \cdots b$, for $a \neq b$.*

Proof: Assume to the contrary that C does contain such a subcycle. Consider the point $p_{a,b}$ of intersection of the curves a and b. There are two cases to consider:

Case (i): a lies below b to the left of $p_{a,b}$. We claim that there is no subsequence $a \cdots b$ in A to the left of p_{ab} (that is, involving elements whose associated intersection points have x-coordinates smaller than that of $p_{a,b}$). Indeed, if such a subsequence exists, then there are curves a' and b' in Γ such that (a, a') and (b, b') are edges in G, a' is above a to the left of $p = a \cap a'$, b' is above b to the left of $q = b \cap b'$, and p lies to the left of q. It is easily seen that in such a case the two edges (a, a'), (b, b') form an anti-diamond in G (see Figure 6), contrary to assumption. Symmetric arguments show that there is no subsequence $b \cdots a$ of A to the right of $p_{a,b}$, no subsequence $b \cdots a$ of B to the left of $p_{a,b}$, and no subsequence $a \cdots b$ of B to the right of $p_{a,b}$.

These arguments imply that A cannot contain a subsequence $a \cdots b \cdots a$, for otherwise A would have to contain either $a \cdots b$ to the left of $p_{a,b}$, or $b \cdots a$ to the right of $p_{a,b}$, both of which are impossible. Similarly, B cannot contain a subsequence $b \cdots a \cdots b$, for that would imply that B would have to contain either $b \cdots a$ to the left of $p_{a,b}$ or $a \cdots b$ to the right of $p_{a,b}$, both of which are impossible.

Hence, if the concatenated sequence C contains an $a \cdots b \cdots a \cdots b$ then, since A cannot contain an $a \cdots b \cdots a$ and B cannot contain a $b \cdots a \cdots b$, the only case to consider is that A contains an $a \cdots b$, where b is (necessarily) to the right of $p_{a,b}$, and B contains an $a \cdots b$, where a is (necessarily) to the left of $p_{a,b}$. In that case, the two intersection points that correspond to the element b of A and to the element a of B in the above subsequence form an anti-diamond (see Figure 7), contradicting our assumption that G is anti-diamond free.

Case (ii): b lies below a to the left of $p_{a,b}$. There are three subcases to consider.

In the first subcase, A contains an $a \cdots b \cdots a$ and B contains b. Reversing the roles of a, b in the analysis of Case (i), we conclude that A does not contain $b \cdots a$ to the left of $p_{a,b}$, and $a \cdots b$ to its right. Hence, there is an intersection point of G labeled a in A to the left of $p_{a,b}$, and another such point labeled a in A to the right of $p_{a,b}$. It is easily verified that the edge (intersection point) e of G labeled by b in B and that edge labeled by a in A that lies on the side of $p_{a,b}$ opposite to e form an anti-diamond (see Figure 8), a contradiction that rules out this subcase.

A symmetric argument excludes the case where A contains a single a and B contains $b \cdots a \cdots b$.

The third possible case is that A contains an $a \cdots b$ and B also contains an $a \cdots b$. But again the first a (that belongs to A) must be to the left of $p_{a,b}$ and the second b (that belongs to B) must be to the right of $p_{a,b}$ and in that case those two are labels of edges that form an anti-diamond; see Figure 9 for an illustration. This completes the proof of the lemma. $\square$

A *run* in C is a maximal contiguous subsequence of identically labeled elements. If we replace each run by a single element, the resulting sequence C^* is a Davenport-Schinzel cycle of order 2 on n symbols, as follows from Lemma 5.1. Hence, the length of C^* is at most $2n - 2$ [21].

Note that it is impossible to have an index $1 \leq$

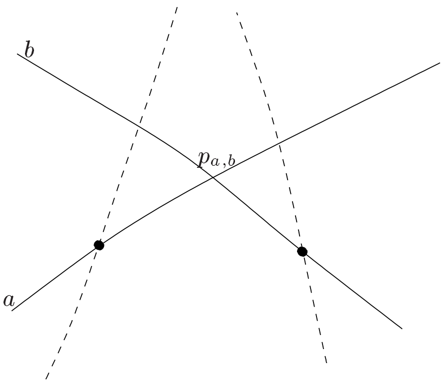

Figure 7: The anti-diamond arising in the second part of Case (i).

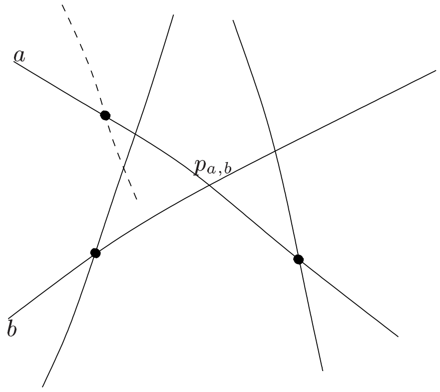

Figure 8: The anti-diamond arising in the first subcase of Case (ii).

$i \leq |G|$ such that the i-th element of A is equal to the $(i+1)(mod|G|)$-st element of A and the i-th element of B is equal to the $(i+1)(mod|G|)$-st element of B. Indeed, if these elements are a and b, respectively, then we obtain two vertices of $\mathcal{A}(\Gamma)$ (the one encoded by the i-th elements of A and B and the one encoded by the $(i+1)$-st elements) that are incident to both a and b, which is impossible. In other words, for each $i = 1, \ldots, |G|$, a new run must begin either after the i-th element of A or after the i-th element of B (or after both). This imply that the length of C^* is greater or equal to $|G|$. Hence we have: $|G| \leq |C^*| \leq 2n - 2$. This completes the proof of Theorem 1.11. $\square$

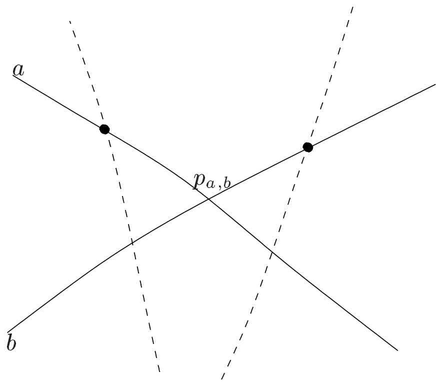

Figure 9: The anti-diamond arising in the third subcase of Case (ii).

6 Pseudolines and Thrackles

Let G be a thrackle with n vertices, whose edges are extendible pseudo-segments. We transform G, using the pseudoline duality, to an intersection graph in an arrangement of a set Γ of n pseudolines. The edge set of G is mapped to a subset E of vertices of $\mathcal{A}(\Gamma)$, with the property that every pair of vertices of E, not sharing a common pseudoline, form a diamond.

Theorem 6.1 *$|E| \leq n$.*

Proof: The proof is an extension, to the case of pseudoline graphs (or, rather, generalized geometric graphs drawn with extendible pseudo-segments), of the beautiful and simple proof of Perles, as reviewed, e.g., in [20].

Fix a pseudoline $\gamma \in \Gamma$ and consider the vertices in $E \cap \gamma$. We say that $v \in E \cap \gamma$ is a *right-turn* (resp., *left-turn*) vertex with respect to γ if, to the left of v, γ lies above (resp., below) the other pseudoline incident to v.

If γ contains three vertices $v_1, v_2, v_3 \in E$, appearing in this left-to-right order along γ, such that v_1 and v_3 are right-turn vertices and v_2 is a left-turn vertex, then all vertices of E must lie on γ, because the intersection of the three (open) double wegdes of v_1, v_2, v_3 is empty, as is easily checked. In this case $|E| \leq n - 1$ and the theorem follows. A similar argument holds when v_1 and v_3 are left-turn and v_2 is a right-turn vertex.

Hence we may assume that, for each $\gamma \in \Gamma$, the left-turn vertices of $E \cap \gamma$ are separated from the right-turn vertices of $E \cap \gamma$ along γ.

For each $\gamma \in \Gamma$, we delete one vertex of $E \cap \gamma$, as follows. If $E \cap \gamma$ consists only of left-turn vertices, or only of right-turn vertices, we delete the rightmost vertex of $E \cap \gamma$. Otherwise, these two groups of vertices are separated along γ, and we delete the rightmost vertex of the left group.

We claim that after all these deletions, E is empty. To see this, suppose to the contrary that there remains a vertex $v \in E$, incident to two pseudolines $\gamma_1, \gamma_2 \in \Gamma$, such that γ_1 lies below γ_2 to the left of v. Clearly, v is a left-turn vertex with respect to γ_1, and a right-turn vertex with respect to γ_2.

The deletion rule implies that, initially, $E \cap \gamma_1$ contained either a left-turn vertex v_1^- that lies to the left of v, or a right-turn vertex v_1^+ that lies to the right of v. Similarly, $E \cap \gamma_2$ contained either a right-turn vertex v_2^- that lies to the left of v, or a left-turn vertex v_2^+ that lies to the right of v. It is now easy to check (see Figure 10) that, in each of the four possible cases, the respective pair of vertices, (v_1^-, v_2^-), (v_1^+, v_2^-), (v_1^-, v_2^+), or (v_1^+, v_2^+), do not form a diamond, a contradiction that shows that, after the deletions, E is empty. Since we delete at most one vertex from each pseudoline, it follows that $|E| \leq n$. $\square$

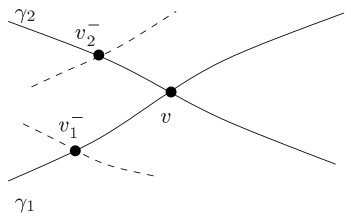

Figure 10: A vertex v remaining after the deletions implies the existence of a pair of vertices that do not form a diamond.

Acknowledgments

The authors would like to thank Pankaj Agarwal, Boris Aronov, János Pach, and Pavel Valtr for helpful discussions concerning the problems studied in this paper.

References

[1] P.K. Agarwal, B. Aronov, T. M. Chan and M. Sharir On levels in arrangements of lines, segments, planes, and triangles *Discrete Comput. Geom.*, 19 (1998), 315–331.

[2] P.K. Agarwal, B. Aronov, J, Pach, R. Pollack and M. Sharir, Quasi-planar graphs have a linear number of edges, *Combinatorica* 17 (1997), 1–9.

[3] P.K. Agarwal, B. Aronov and M. Sharir, On the complexity of many faces in arrangements of pseudo-segments and of circles, in preparation.

[4] P.K. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir and S. Smorodinsky, Lenses in arrangements of pseudodisks and their applications, in preparation.

[5] P.K. Agarwal and M. Sharir, Pseudoline arrangements: Duality. algorithms and applications, manuscript, 2001.

[6] M. Ajtai, V. Chvátal, M. Newborn and E. Szemerédi, Crossing-free subgraphs, *Ann. Discrete Math* 12 (1982), 9–12.

[7] G. Cairns and Y. Nikolayevsky, Bounds for generalized thrackles, *Discrete Comput. Geom.*, 23 (2000), 191–206.

[8] T.M. Chan, On levels in arrangements of curves, *Proc. 41st IEEE Symp. Found. Comput. Sci.* (2000), 219–227.

[9] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete Comput. Geom.* 5 (1990), 99–160.

[10] K. Clarkson and P. W. Shor Applications of random sampling in computational geometry, II *Discrete Comput. Geom.* 4 (1989), 387–421.

[11] T. K. Dey, Improved bounds on planar k-sets and related problems, *Discrete Comput. Geom.* 19 (1998), 373–382.

[12] T. Dey and J. Pach, Extremal problems for geometric hypergraphs, *Discrete Comput. Geom.* 19 (1998), 473–484.

[13] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-verlag, Heidelberg, 1987.

[14] I. Fáry, On straight-line representation of planar graphs, *Acta Sciientiarum Mathematicarum (Szeged)* 11 (1948), 229–233.

[15] M. Katchalski and H. Last, On geometric graphs with no two edges in convex position, *Discrete Comput. Geom.* 19 (1998), 399–404.

[16] F. T. Leighton, *Complexity Issues in VLSI*, MIT Press, Cambridge, MA, 1983.

[17] F. Levi, Die Teilung der projektiven Ebene durch Gerade oder Pseudogerade, *Ber. Math-Phys. Kl. Sächs. Akad. Wiss.* 78 (1926), 256–267.

[18] L. Lovász, J, Pach and M. Szegedy, On Conway's thrackle conjecture, *Discrete Comput. Geom.* 18 (1997), 369–376.

[19] J. S. B. Mitchell and J. O'Rourke. Computational geometry column 42. *Internat. J. Comput. Geom. Appl.* (2001). Also in SIGACT News 32(3):63-72 (2001), Issue 120. See also: http://www.cs.smith.edu/ orourke/TOPP/

[20] J. Pach, Geometric graph theory, in *Surveys in Combinatorics* (J.D. Lamb and D.A. Preece, eds.), London Mathematical Society Lecture Note Series 267, Cambridge University Press, 1999, 167–200.

[21] M. Sharir and P.K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

[22] L. Székely, Crossing numbers and hard Erdős problems in discrete geometry, *Combinatorics, Probability and Computing* 6 (1997), 353–358.

[23] H. Tamaki and T. Tokuyama, A characterization of planar graphs by pseudo-line arrangements, *Proc. 8th Annu. Internat. Sympos. Algorithms Comput.* (ISAAC '97), Springer-Verlag Lecture Notes Comput. Sci., Vol. 1350, 1997, 133–142.

[24] G. Tóth Point sets with many k-sets *Discrete Comput. Geom.* 26 (2001), 187–194.

[25] W. T. Tutte, Toward a theory of crossing numbers, *J. Combinat. Theory* 8 (1970), 45–53.

[26] P. Valtr, Graph drawings with no k pairwise crossing edges, *Lecture Notes Comput. Sci. Springer-Verlag* 1353 (1997), 205–218.

[27] P. Valtr, On geometric graphs with no k pairwise parallel edges, *Discrete Comput. Geom.* 19 (1998), 461–469.

[28] P. Valtr, Generalizations of Davenport-Schinzel sequences, in *Contemporary Trends in Discrete Mathematics* (J. Nešetřil, J. Kratochvil, F.S. Roberts, R.L. Graham, Eds.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 49 (1999), 349–389.