

How to Cut Pseudoparabolas into Segments

Hisao Tamaki and Takeshi Tokuyama

IBM Research Division, Tokyo Research Laboratory,
1623-14 Shimo-tsuruma, Yamato, Kanagawa 242, Japan

Abstract. Let Γ be a collection of unbounded x -monotone Jordan arcs intersecting at most twice each other, which we call pseudoparabolas, since two axis parallel parabolas intersect at most twice. We investigate how to cut pseudoparabolas into the minimum number of curve segments so that each pair of segments intersect at most once. We give an $\Omega(n^{4/3})$ lower bound and $O(n^{5/3})$ upper bound on the number of cuts. We give the same bounds for an arrangement of circles. Applying the upper bound, we give an $O(n^{23/12})$ bound on the complexity of a level in an arrangement of pseudoparabolas, and an $O(n^{11/6})$ bound on the complexity of a combinatorially concave chain of pseudoparabolas. We also give some upper bounds on the number of transitions of the minimum weight matroid base when the weight of each element changes as a quadratic function of a single parameter.

1. Introduction

Arrangements of curves in a plane is a major research target in computational geometry. Combinatorial complexities of parts of arrangements such as a cell, many cells, k -levels, $\leq k$ -levels, and x -monotone chains play key roles in designing algorithms on geometric optimization and motion planning problems [4], [7], [14], [15].

Although arrangements of lines and line segments are most popular, an arrangement of curves that satisfy the condition that each pair of curves intersect at most s times for a given constant s , is also important in both theory and applications [11], [16]. If each pair of curves intersects at most once, the arrangement is called an arrangement of pseudolines, to which many results on an arrangement of lines generalizes. For example, the complexity of the k -level of an arrangement of n pseudolines is known to be $O(\sqrt{kn})$ [19], [4].

We focus on the case $s = 2$ in this paper. A familiar example of such an arrangement is that of axis-parallel parabolas (Fig. 1), in which two curves intersect at most twice. An arrangement of axis-parallel parabolas is used in dynamic computational geometry, since it shows the transition of the list of pairs of points among a set of linearly moving




Fig. 1. Arrangement of parabolas.

points sorted with respect to their L_2 distances. The complexity of the lower envelope and the k -level of an arrangement of parabolas gives the number of combinatorial changes on the nearest pair and the k th nearest pair, respectively [1]. Also, the complexity of topological change on the L_2 minimum spanning tree can be formulated into a problem on an arrangement of parabolas [10].

More generally, we consider an arrangement of unbounded x -monotone Jordan curves intersecting each other at most twice. Such an arrangement is called an arrangement of 2-intersecting curves in the literature; however, for convenience's sake, we call it an arrangement of *pseudoparabolas*.

It is often more difficult to analyze the complexity of an arrangement of curves than an arrangement of lines or pseudolines. For example, the only upper bound previously known on the k -level complexity of an arrangement of parabolas is $O(kn)$, which is the same as the bound for $\leq k$ -levels [16].

The aim of this paper is to link the complexity of an arrangement of pseudoparabolas to that of an arrangement of pseudolines. Our approach is to split pseudoparabolas by cut points, generating an arrangement of *pseudosegments* in which each pair of pseudo-segments intersect at most once. For example, the arrangements of Fig. 1 can be converted into an arrangement of pseudo-segments by introducing seven cuts (Fig. 2). We call the minimum number of cuts required to make an arrangement Γ into an arrangement of pseudosegments the *cutting number* of Γ .




Fig. 2. Cut points.

Our main results are the following two theorems:

Theorem 1.1. *There exists an arrangement of axis parallel parabolas whose cutting number is $\Omega(n^{4/3})$.*

Theorem 1.2. *An arrangement of pseudoparabolas can be cut into an arrangement of pseudosegments with $O(n^{5/3})$ cuts.*

We also give the same bounds for the cutting number of an arrangement of circles.

The lower bound of Theorem 1.1 is derived from the lower-bound example [4], [5] on the complexity of n cells in an arrangement of lines.

The upper bound of Theorem 1.2 is derived from an inequality of Lovász's used in the proof of his fractional-covering theorem [13], combined with extremal graph theory [2] and a probabilistic method [3], [16]. A greedy algorithm outputs cuts attaining this upper bound.

Combining Theorem 1.2 with the known upper bound on the level complexity of an arrangement of pseudolines, we derive a nontrivial $O(n^{23/12})$ upper bound on the complexity of a level of an arrangement of pseudoparabolas. The technique used here is such that any improved upper bound for pseudolines will lead to an improved upper bound for pseudoparabolas. Thus, Theorem 1.2 establishes an important link between the complexities of arrangements of these two types.

We also give some upper bounds on the number of transitions of the minimum weight matroid base when the weight of each element changes as a quadratic function of a single parameter t .

2. Preliminaries

Let Γ be an arrangement of pseudoparabolas. The arrangement subdivides the plane into faces. We use the terms *cell*, *edge*, and *vertex* for two-, one-, and zero-dimensional faces, respectively. When two pseudoparabolas intersect twice, they form a closed curve consisting of a segment of each pseudoparabola: we call this closed curve a *lens* formed by these two pseudoparabolas.

The boundaries of the shaded regions in Fig. 3 are lenses. We say a lens is a 1-lens if no curve crosses the lens. Consequently, a 1-lens consists of exactly two edges of the arrangement. There exist no 1-lens in Fig. 3.

We define a hypergraph $H(\Gamma)$, whose node set is the set of edges of the arrangement Γ . A set of nodes of $H(\Gamma)$ forms a hyperedge if and only if its corresponding set of edges of the arrangement forms a lens.

A node covering of a hypergraph H is a subset of the node set of H such that every hyperedge contains at least one node of the set.

A node covering with minimum size (number of nodes) is called a minimum covering. The size of a minimum covering is called the covering number.

The following is a key lemma for our upper bound of the cutting number:

Lemma 2.1. *The cutting number of Γ is equal to the covering number of $H(\Gamma)$.*




Fig. 3. Lenses.

Proof. Given a minimum covering C of $H(\Gamma)$, we cut all edges of Γ associated with nodes in C . Then all lens are cut, that means all pairs of curve segments after the cut intersect at most once. On the contrary, given a minimum cut of Γ , consider the collection of the edges cut. Then the associated node set of $H(\Gamma)$ is a covering. \square

Corollary 2.2. *The cutting number of Γ is not less than the number of 1-lenses.*

3. Lower Bound

We give a proof of Theorem 1.1 in this section. We use the following well-known fact (see [4] and [5]):

Theorem 3.1. *There exists a highly degenerate arrangement \mathcal{A} of lines which has a set V of n vertices with total degree $\Omega(n^{4/3})$.*

Due to the construction given in [4], we can further assume that each line of the arrangement \mathcal{A} has a positive slope which is not larger than a given constant, say 0.5.

Proposition 3.2. *There exists an arrangement of $O(n)$ parabolas whose cutting number is $\Omega(n^{4/3})$.*

Proof. Let \mathcal{A} be the line arrangement introduced above. We consider an axis-parallel parabola γ , defined by $y = cx^2$ for a positive constant c , and draw a copy $\gamma(v)$ which has its peak at v for each v in V . Hence, we have n parabolas in total.

Now, we translate each line, so that the degeneracy of \mathcal{A} is resolved, and each line through a point $v = (x_v, y_v)$ in V is translated so that it is tangent to $\gamma(v)$. More precisely, each line with a slope a is translated with a vector $(a/2c, a^2/4c)$, so that the line $y - y_v = a(x - x_v)$ is translated to $y - y_v + a^2/4c = a(x - x_v)$, which is tangent to $\gamma(v)$. Note that the translation vector is independent of v .

We consider the union of the above perturbation of \mathcal{A} and the parabolas. Then the number of vertices in this arrangement at which two curves are tangent is $\Omega(n^{4/3})$. If we replace $\gamma(v)$ with its translated copy with the (infinitesimal) translation vector $(0, -\varepsilon)$ for a very small positive ε , the arrangement has $\Omega(n^{4/3})$ 1-lenses. The arrangement has n lines, as well as n parabolas. However, we can replace each curve $y = f(x)$ with $y = f(x) + x^2$ to convert into an arrangement of $2n$ parabolas with the same combinatorial property. \square

4. Upper Bound

It suffices to give an upper bound for the covering number of hypergraph $H(\Gamma)$. We recall some notations on hypergraphs. The degree $d(x)$ of a node x of a hypergraph is the number of hyperedges containing x . The maximum degree in a hypergraph H is denoted by $d(H)$.

A *simple k -matching* of H is a set \mathcal{M} of hyperedges such that each node belongs to at most k of them. The maximum number of hyperedges in a simple k -matching is denoted by $v_k(H)$ (this is denoted by $\tilde{v}_k(H)$ in [13]). Note that $v_{d(H)}(H)$ is the number of hyperedges in H . We remove the argument H from functions d and v if no confusion arises. A greedy algorithm for computing a covering is the following:

1. Find a node of maximum degree;
2. Insert the node into the covering, and remove the node and all hyperedges containing it from H ;
3. If all hyperedges are covered, Exit; Else GOTO 1;

Lovász [13] shows that the greedy algorithm achieves a covering of size at most $\log d(H) + 1$ times the size of the covering of H . The following is the key inequality in his proof. Let t be the size of the covering of H obtained by the greedy algorithm. Then

$$t \leq \frac{v_1}{1 \cdot 2} + \frac{v_2}{2 \cdot 3} + \dots + \frac{v_{d-1}}{(d-1) \cdot d} + \frac{v_d}{d}. \tag{1}$$

Consequently, the minimum covering number of H is also bounded by the right-hand side of (1). Therefore, we want to estimate $v_k(H(\Gamma))$ for $k = 1, 2, \dots, d$.

Suppose we have a simple 1-matching \mathcal{M} of $H(\Gamma)$ of size M . Recall that a hyperedge in $H(\Gamma)$ is a lens in Γ .

We define a bipartite graph $G(\mathcal{M})$. The vertex set is $S_1 \cup S_2$, where S_1 and S_2 are disjoint, $|S_1| = |S_2| = |\Gamma| = n$, with associated bijections $\gamma_1: S_1 \rightarrow \Gamma$ and $\gamma_2: S_2 \rightarrow \Gamma$.

We draw an edge between a node u of S_1 and v of S_2 if and only if the associated curves $\gamma_1(u)$ and $\gamma_2(v)$ form a lens which is associated with a hyperedge in \mathcal{M} , and $\gamma_1(u)$ is above $\gamma_2(v)$ within the lens. Here, a curve γ is *above* another curve μ within their lens L if the vertical downward ray from a point on $\gamma \cup L$ intersects μ .

By definition, $G(\mathcal{M})$ has $2n$ vertices. It is clear that the number of edges in $G(\mathcal{M})$ is the size of the matching \mathcal{M} .

First, we bound the size $v_1(H(\Gamma))$ of the maximum simple 1-matching.

Lemma 4.1. *Suppose \mathcal{M} is a simple 1-matching of $H(\Gamma)$. Then $G(\mathcal{M})$ does not contain $K_{3,4}$ as a subgraph.*

Proof. Assume that $G(\mathcal{M})$ contains a copy of $K_{3,4}$. Then we have three curves C_1, C_2, C_3 and four curves D_1, D_2, D_3, D_4 such that each pair (C_i, D_j) makes a lens, in which C_i is above D_j , for $1 \leq i \leq 3, 1 \leq j \leq 4$. Furthermore, because \mathcal{M} is a 1-matching, no edge in the arrangement constructed from these seven curves is contained in more than one such lens. Let $A(C)$ denote the arrangement of $\{C_1, C_2, C_3\}$ and $A(D)$ that of $\{D_1, D_2, D_3, D_4\}$. If an edge e of the arrangement is located on D_1 , and e is below two curves C_1 and C_2 , e must be on both lenses (C_1, D_1) and (C_2, D_1) . This means that the two arrangements $A(C)$ and $A(D)$ intersect each other only at points that are on the upper envelope of $A(C)$. Similarly, those intersection points must also be on the lower envelope of $A(D)$. Since there must be 12 lenses, the number of intersections must be at least 24. However, the upper envelope of $A(C)$ has at most 5 edges, and the lower envelope of $A(D)$ has at most 7 edges [1]. Because each pair of curves intersects at most twice, the number of intersecting points cannot exceed 22, which is a contradiction. \square

Remark. With a more careful analysis, we can also show that $G(\mathcal{M})$ does not contain $K_{3,3}$.

We use the following result in extremal graph theory, which can be found in [2] in a more general form (p. 73, Lemma 7).

Lemma 4.2. *Let G be a bipartite graph, with n vertices on each side, which does not contain $K_{s,t}$ as a subgraph. Suppose G contains $m = yn$ edges. Then $n \binom{y}{s} \leq (t-1) \binom{n}{s}$.*

Theorem 4.3. *$G(\mathcal{M})$ contains $O(n^{5/3})$ edges. Hence, $v_1(H(\Gamma)) = O(n^{5/3})$.*

Proof. We substitute $s = 3$ and $t = 4$ in the above lemma, and obtain $y = O(n^{2/3})$. \square

Next, we show that $v_k(H(\Gamma)) = O(n^{5/3}k^{1/3})$ for general k . In the following, we assume $2 \leq k \leq n/7$: the case $k = 1$ is already dealt with and the case $k > n/7$ is trivial. We apply the probabilistic argument similar to the one Sharir [16] used for analyzing the complexity of the $\leq k$ level of curves.

Suppose we have a simple k -matching \mathcal{M} , and associated set $\mathcal{L}(\mathcal{M})$ of lenses. Assume \mathcal{M} has v_k hyperedges (i.e., $\mathcal{L}(\mathcal{M})$ has v_k lenses).

For each lens L bounded by two curves C_1 and C_2 , an arrangement edge on L is called *extremal* if it contains one of the intersection points of C_1 and C_2 . Obviously, there are at most four extremal edges associated with L .

Now, we choose a sample Y of $r = \lfloor n/k \rfloor$ curves from Γ . We say a lens L is a *near 1-lens* in the sample if (1) L is a lens consisting of two curves of Y and (2) for each arrangement edge e of Γ that is extremal in L , L is the only lens which consists of two curves of Y and contains e (not necessary as an extremal edge).

We consider the set of near 1-lenses of the sample Y , and consider the associated set of hyperedges \mathcal{M}_0 of $H(Y)$.

Lemma 4.4. \mathcal{M}_0 is a simple 1-matching of $H(Y)$.

Proof. Suppose an edge e of the arrangement of Y is contained in two near 1-lenses L and L' . Let C be the curve on which e is located. Then L and L' contain intervals I and I' of C . Both intervals must contain e ; thus, at least one endpoint of either I or I' must be contained in the other interval. Without loss of generality, we assume an endpoint of I is contained in I' . This means that an extremal edge of L in the original arrangement is contained in L' , which contradicts the definition of a near 1-lens. \square

Now, analyze the expected number of near 1-lenses in the random sample Y . A lens L of $\mathcal{L}(\mathcal{M})$ becomes a near 1-lens if (1) both of its bounding curves are in the sample and (2) no other curve contributing a lens containing an extremal edge of L is in the sample. Since \mathcal{M} is a k -matching, the number of such curves in (2) is at most k for each extremal edge. For a fixed L , the number of samples Y of size r satisfying these conditions is at least $\binom{n - 4k - 2}{r - 2}$, so the probability that the lens L becomes a near 1-lens is at least $\binom{n - 4k - 2}{r - 2} / \binom{n}{r}$. By a standard calculation (see [3] and [16]), we have

$$\frac{\binom{n - 4k - 2}{r - 2}}{\binom{n}{r}} \geq \frac{r(r - 1)}{n(n - 1)} \left(1 - \frac{r - 2}{n - 4k - 1}\right)^{4k}. \tag{2}$$

Since $r = \lfloor n/k \rfloor$ and we are assuming $k \leq n/7$, we have $(r - 2)/(n - 4k - 1) \leq 2/k$ and $r(r - 1)/n(n - 1) \geq k^{-2}/2$ and hence the right-hand side of (2) is at least $\frac{1}{2}k^{-2}(1 - (2/k))^{4k} = \Omega(k^{-2})$. Therefore, the expected number of near 1-lenses in Y is $\Omega(k^{-2}v_k)$. On the other hand, the bound on the size of a simple 1-matching in the sample implies that the number of near 1-lenses in Y must be $O(r^{5/3})$. Solving $\Omega(k^{-2}v_k) \leq O((n/k)^{5/3})$ for v_k , we have:

Lemma 4.5. $v_k = O(n^{5/3}k^{1/3})$.

Now, we compute the right-hand side of Lovász's inequality (1):

$$\begin{aligned} \frac{v_1}{1 \cdot 2} + \frac{v_2}{2 \cdot 3} + \dots + \frac{v_{d-1}}{(d-1) \cdot d} + \frac{v_d}{d} &= O\left(n^{5/3} \sum_{k=1}^n \frac{k^{1/3}}{k(k+1)}\right) \\ &= O(n^{5/3}). \end{aligned}$$

This proves Theorem 1.2.

5. Applications

Level Complexity

Let Γ be an arrangement of n pseudoparabolas. The *level* of an edge e of the arrangement Γ is the number of edges that intersect with the vertical half-line downward from an internal point on e . This notion is well-defined since the above number is independent of the choice of any internal point of e . It is well-known that the closure of the union of all edges with a given level k is a connected curve and separates the plane. This curve is called the k -level of Γ . The complexity of the k -level of Γ is the number of edges whose level is k .

Theorem 5.1. *The complexity of the k -level of Γ is $O(n^{23/12})$.*

Proof. Without loss of generality, we assume that the arrangement is simple, that is, no three curves intersect at a point. We consider the set P of cutting points that cut the pseudoparabolas of Γ into segments. We know that $m = |P| = O(n^{5/3})$. We subdivide the plane into $m + 1$ slabs with m vertical lines through points of P . Inside a slab, the arrangement can be considered as an arrangement of pseudolines.

Let X_i be the number of vertices of the arrangement located in the i th slab S_i . Then $\sum_{i=1}^{m+1} X_i = O(n^2)$.

Suppose that exactly n_i curves contribute to the k -level in S_i . The k -level of Γ inside S_i is a level of these n_i curves.

Suppose a curve γ is on the k' -level at the left end of the slab S_i , and contributes to the k -level in S_i . Then γ must have at least $|k - k'|$ vertices on itself in S_i . Therefore, the arrangement must have at least $n_i^2/2$ vertices in S_i . This means that $n_i = O(\sqrt{X_i})$.

The complexity of a level of n_i pseudolines is $O(n_i \sqrt{n_i})$ [4], which is $O(X_i^{3/4})$.

Thus, the complexity of the k -level of Γ is $O(\sum_{i=1}^{m+1} X_i^{3/4}) = O(m(n^2/m)^{3/4})$. Since $m = O(n^{5/3})$, we obtain $O(n^{23/12})$. □

Note that this result improves the known bound of $O(kn)$ [16] when $k > n^{11/12}$. Also note that the bound would be automatically improved further, if we had a better bound either on the level complexity of an arrangement of pseudolines or on the cutting number of pseudoparabolas.

Complexity of Combinatorial Concave Chains

Given an arrangement \mathcal{F} of x -monotone curves, a *combinatorial concave chain* (cc-chain, for short) c is an x -monotone chain in the arrangement \mathcal{F} satisfying the following concave condition: If two curves γ_1 and γ_2 intersects at a point p and the chain c lies on γ_1 to the left of p and on γ_2 to the right of p , respectively, then γ_1 stabs γ_2 from below at p (curves are directed from left to right). See Fig. 4.

Given an arrangement Γ of pseudoparabolas and a set \mathcal{C} of cc-chains of that arrangement, a *chain vertex* of \mathcal{C} is a vertex of Γ at which a cc-chain of \mathcal{C} switches from one



Fig. 4. Combinatorial concave chain.

curve of Γ to another; a *chain edge* of C is the segment of a curve between two consecutive chain vertices of a cc-chain of C . To avoid confusion, we refer to an edge of the given arrangement as an *arrangement edge*. We define the *length* of a cc-chain to be the number of chain edges of the cc-chain. We also refer to the length of a cc-chain as the *complexity* of the cc-chain.

Theorem 5.2. *The worst-case complexity of a cc-chain in an arrangement of n pseudoparabolas is $\Omega(n^{4/3})$ and $O(n^{11/6})$.*

Proof. The lower bound can be obtained from the degenerate arrangement which was used in Proposition 3.2. This arrangement has $O(n)$ lines and $\Omega(n)$ “heavy” nodes to each of which $\Omega(n^{1/3})$ lines are incident. Moreover, the heavy vertices are lattice points in an $\sqrt{n} \times \sqrt{n}$ integral lattice of the plane (see p. 112 of [4]). We rotate the arrangement clockwise with a suitable small angle, so that the vertical column of the lattice is transformed into a line with a positive slope. Each “heavy” vertex is replaced by a tiny concave chain by using a suitable perturbation. We connect the rightmost vertex of this tiny chain with the leftmost vertex of the tiny chain of the next (from top to bottom) lattice vertex in the same column with a very steep parabola. Moreover, we connect the topmost lattice point in a column with the bottom lattice point in the next column using a very sharp parabola, so that the parabola joins two concave chains into one concave chain. See Fig. 5. Then this arrangement has a combinatorial concave chain of length $\Omega(n^{4/3})$.

The upper bound can be obtained by subdividing the plane into $m + 1 = O(n^{5/3})$ slabs given in the proof of Theorem 5.1. If there are n_i segments in the i th slab, the complexity of the cc-chain within the slab is $O(n_i)$. Since $\sum_{i=1}^{m+1} n_i^2 = O(n^2)$, $\sum_{i=1}^{m+1} n_i = O(n\sqrt{m}) = O(n^{11/6})$. \square

Theorem 5.3. *The worst-case complexity of a cc-chain of n curves where each pair of curves intersects at most three times is $\Omega(n^2/\log n)$.*

Proof. We can make any x -monotone chain in a pseudoline arrangement into a combinatorial concave chain of n 3-intersecting curves by using local changes (Fig. 6). Thus, the theorem follows from Matoušek’s $\Omega(n^2/\log n)$ lower bound on the length of an x -monotone chain [14] in an arrangement of pseudolines. \square




Fig. 5. Construction of a concave chain of size $\Omega(n^{4/3})$.

We call a set of cc-chains of an arrangement *edge-disjoint* if no arrangement edge is shared by more than one cc-chain in the set. We are interested in bounding the total complexity of k edge-disjoint cc-chains.

Theorem 5.4. *Let Γ be an arrangement of n pseudoparabolas. Suppose $k \leq n^{1-\delta}$ for some positive constant $\delta < \frac{6}{11}$. Then the total number of chain edges in any set of k edge-disjoint unbounded cc-chains of Γ is $O(n^{2-\epsilon})$ where $\epsilon = \delta/12$.*

Since the proof of this theorem is rather long, we delegate it to Section 7.

Transitions of Minimum Matroid Base and MST

Let E be a finite set and let \mathcal{B} be a family of subsets of E . The pair (E, \mathcal{B}) is called a *matroid* $M(E, \mathcal{B})$, and the elements of \mathcal{B} are the *bases* of $M(E, \mathcal{B})$, if the following two axioms hold [18]:

- (A1) For any $B, C \subset E$ with $B \neq C$, if $B \in \mathcal{B}$ and $C \subset B$, $C \notin \mathcal{B}$.




Fig. 6. Local change of an x -monotone chain of lines for a concave chain of 3-intersecting curves.

(A2) For any $B, B' \in \mathcal{B}$ with $B \neq B'$ and for any $e \in B - B'$, there exists $e' \in B' - B$ such that $(B - \{e\}) \cup \{e'\} \in \mathcal{B}$.

For instance, let \mathcal{T} be a set of spanning trees in an undirected connected graph $G = (V, E)$; then (E, \mathcal{T}) forms a matroid and \mathcal{T} is a set of bases [18].

The number $|B|$ of elements of a base $B \in \mathcal{B}$ is independent of the choice of B [18], and is denoted by p . Let $m = |E|$, and assume the elements of E to be indexed from 1 through m . We assume that each element i has a real-valued weight $w_i(t)$ that is a function in the parameter t . The minimum (resp. maximum) weight base is the one in which the sum of weights of elements is minimum (resp. maximum).

If the weight functions of two elements have a constant number of intersections, we have an $O(m^2)$ trivial upper bound on the number of transitions of the minimum weight base of $M(E, \mathcal{B})$. If $w_i(t)$ is linear, this was improved to $O(m \min\{\sqrt{p}, \sqrt{m-p}\})$ [8], [9]. Recently, Eppstein [6] constructed a matroid with linear weights for which the number of transitions of the minimum matroid base is $\Omega(m^{2/3} p^{2/3})$.

If the weight functions are quadratic, it is not even clear if the $O(mp)$ upper bound holds. From Theorem 5.4, we have the following bound:

Theorem 5.5. *If the weight functions are quadratic, the number of transitions of the minimum weight base is $O(m^{23/12} p^{1/12})$ if $p > m^{5/11}$, and $O(m^{2-1/22})$ if $p \leq m^{5/11}$.*

Proof. The locus $y = w_i(t)$ is a parabola on the t - y plane for each $i = 1, 2, \dots, m$. Hence, we have an arrangement of m parabolas. We assume without loss of generality that the arrangement of parabolas is simple; otherwise, we can perturb it to obtain a simple arrangement without decreasing the number of transitions.

Let $B_{\min}(t)$ be the minimum weight at t , and let $I_{\min}(t)$ be the set of indices j such that $e_j \in B_{\min}(t)$. Let $Y(t)$ be the trajectory $\{(t, w_j(t)) : j \in I_{\min}(t)\}$ on the t - y plane.

Suppose that a transition occurs at $t = t_0$. Then there must be a pair of bases B and B' such that both attain the minimum weight at t_0 . Moreover, we can assume that there exists an infinitesimally small positive real number ε such that B and B' are the minimum weight bases at $t_0 - \varepsilon$ and $t_0 + \varepsilon$, respectively.

Then, from axiom (A2), there exists e_i and e_j in E such that $B' = (B - \{e_i\}) \cup \{e_j\}$. Since both B and B' are minimum bases, $w_i(t_0) = w_j(t_0)$, and the chain

$$y = \begin{cases} w_i(t) & \text{for } t \leq t_0, \\ w_j(t) & \text{for } t > t_0 \end{cases}$$

is combinatorial concave. Thus, we can observe that $Y(t)$ is a union of p combinatorial concave chains, so that the number of transitions is the sum of complexities of these chains. Note that this is not the complexity of $Y(t)$, since we ignore intersections between concave chains. Hence, the theorem follows from Theorem 5.4. □

Corollary 5.6. *Let G be a graph with m edges and n nodes, and suppose that each edge has a weight function which is quadratic in a parameter t . Then the number of transitions in its minimum spanning tree is $O(m^{23/12} n^{1/12})$. In particular, the number of transitions of the minimum spanning tree of a complete graph K_n is $O(n^{4-1/12})$.*

Remark. For a *trivial matroid*, where \mathcal{B} is the set of all subsets of size p of E , the number of transitions of the minimum weight base equals the complexity of the p -level of the arrangement of parabolas, and, hence, we have an $O(\min\{mp, m^{23/12}\})$ bound by combining Theorem 5.1 with an upper bound of complexity of the at-most- p levels [16].

6. Related Topics

Arrangement of Circles

We are given an arrangement of n circles. Although a pair of circles intersects at most twice, a circle is not an x -monotone curve. However, we can cut each circle with its horizontal diameter, and divide it into an upper half-circle and a lower half-circle. We can connect two vertical downward (resp. upward) rays to an upper (resp. lower) half-circle at its two endpoints, and obtain an x -monotone curve separating the plane. It is easy to see that every pair of curves intersects at most twice.

Thus, we have a family of pseudoparabolas. We can now apply our upper bound results in the previous section. Also, the lower bound example for the cutting number of pseudoparabolas in Section 3 can be easily adapted to an arrangement of circles. Indeed, we place a unit circle at each heavy vertex of the arrangement \mathcal{A} in Proposition 3.2, translate lines and circles so that $O(n^{4/3})$ 1-lenses are created. We can replace lines with a circle with a very large diameter without changing the combinatorial structure.

Thus, we have the following:

Theorem 6.1. *Using $O(n^{5/3})$ cuts, an arrangement of circles can be transformed to an arrangement of pseudosegments. There exists an example for which $\Omega(n^{4/3})$ cuts are required.*

Cutting a $2t$ Intersecting Family

It is desired to extend the upper bound on a cutting number to that for an arrangement of curves such that each pair of curves intersect at most t times. Unfortunately, in the worst case, we need to cut at $\Omega(n^2)$ points in order to make an arrangement of 3-intersecting curves into an arrangement of 2-intersecting curve segments. The lower-bound example is obtained as follows: Consider an arrangement of n lines. At each vertex v of the arrangement, we locally replace one of the lines meeting there with a curve intersecting the other line three times to make a pair of two 1-lenses as in Fig. 6. In order to cut the curve so that each pair of curve segments intersect at most twice, we must cut one of the pair of lenses for each v ; hence, we need $\Omega(n^2)$ cuts.

More generally, $\Omega(n^2)$ cuts are needed to make $(2t + 1)$ -intersecting family to $2t$ -intersecting family in the worst case. On the other hand, for a family with even intersecting-numbers, we have the following result:

Theorem 6.2. *Given an arrangement of curves in which every pair of curves intersects at most $2t$ times, we can cut it at $O(n^{2-1/\beta(t)})$ points to make it an arrangement of curve*

segments, in which every pair of segments intersects at most $2t - 1$ times. Here, $\beta(t)$ is the minimum positive number y satisfying $y^2 \geq 4t\lambda_{2t}(y)$, where $\lambda_{2t}(y)$ is the length of a Davenport–Schinzel sequence of degree $2t$ on y characters.

Proof. Consider two curves P and P' such that P is above P' at $x = -\infty$. We consider the region $R(P, P')$ of all points which is below P' and above P . Then $R(P, P')$ consists of t connected components. The $2t$ -lens $L(P, P')$ is the set of arrangement edges bounding $R(P, P')$. It suffices to cut at least one edge in every $2t$ -lens. Now, we define \mathcal{M} and $G(\mathcal{M})$ analogously to Section 3 for the hypergraph associated with the set of $2t$ -lenses. The only difference from the case where $t = 1$ is the analogue of Lemma 4.1, in which we can prove that $G(\mathcal{M})$ does not contain $K_{\beta(t), \beta(t)}$. This implies the theorem. \square

7. Proof of Theorem 5.4

Let $\delta < \frac{6}{11}$ be a positive constant and let $\varepsilon = \delta/12$. Suppose $k \leq c_0 n^{1-\delta}$ for a sufficiently small constant c_0 defined later. Let Γ be an arrangement of n pseudoparabolas and let \mathcal{C} be a set of k edge-disjoint unbounded cc-chains in Γ . Let $\mathcal{L}(\Gamma)$ denote the set of lenses of Γ and $E(\mathcal{C})$ the set of chain edges of \mathcal{C} . Our goal in this section is to prove $|E(\mathcal{C})| \leq 7n^{2-\varepsilon}$. The proof is by contradiction: we suppose $|E(\mathcal{C})| > 7n^{2-\varepsilon}$ and devote the whole section to derive a contradiction.

In this section, for brevity, we often refer to pseudoparabolas of Γ simply as “parabolas.” We use the following notation and definitions. We say a parabola P is *steeper than* a parabola Q if and only if P is above Q at the left-infinity. If two parabolas P and Q of Γ , P steeper than Q , intersect at two points and thus form a lens, we denote this lens by $\text{lens}(P, Q)$; note that P forms the lower boundary of this lens. For each $e \in E(\mathcal{C})$, we denote by $\text{parab}(e)$ the parabola of Γ of which e is a segment. Similarly, we denote by $\text{chain}(e)$ the cc-chain of \mathcal{C} that contains e . For each chain edge or arrangement edge e , we denote by $\text{left}(e)$ (resp. $\text{right}(e)$) the left (resp. right) endvertex of e . We also write $\text{left}(L)$ and $\text{right}(L)$ to denote the left and right endvertices of lens L . For two points u and v in the plane such that u lies to the left of v (i.e., the x -coordinate of u is smaller than that of v), we denote by $\text{slab}[u, v]$ the vertical closed slab whose left boundary contains u and whose right boundary contains v . We denote by $\text{slab}(u, v]$, $\text{slab}[u, v)$, or $\text{slab}(u, v)$ the similar slab which is open at the left, right, or both boundaries, respectively. We also use the notation $\text{slab}[u, \infty)$, $\text{slab}(u, \infty)$, $\text{slab}(-\infty, u]$, and $\text{slab}(-\infty, u)$ to denote slabs unbounded in the specified direction. Moreover, when the left end of the slab is specified by the *right* endvertex of an edge or lens X , we abbreviate $\text{slab}(\text{right}(X), \cdot)$ to $\text{slab}(X, \cdot)$. Similarly we abbreviate $\text{slab}(\cdot, \text{left}(X))$ to $\text{slab}(\cdot, X)$, and use the same convention for closed slabs.

Let $E \subseteq E(\mathcal{C})$ be a set of chain edges of \mathcal{C} . An *adjacent pair* of E is an ordered pair $(e, f) \in E \times E$ such that e and f are segments of the same parabola, e lies to the left of f , and $\text{parab}(e) = \text{parab}(f)$ does not contain any chain edge of E within $\text{slab}(e, f)$. We call E *fully wedged* if, for every adjacent pair (e, f) of E , there is a parabola of Γ that is steeper than $\text{parab}(e)$ and forms a lens with $\text{parab}(e)$ completely contained in $\text{slab}(e, f)$. We call a set of lenses $\mathcal{L} \subset \mathcal{L}(\Gamma)$ *upper-disjoint* if, for any two distinct lenses

of \mathcal{L} whose upper boundaries are defined by the same parabola, their upper boundaries are disjoint from each other. Given a set of chain edges E , the *weight* of a lens L with respect to E , denoted by $\text{weight}_E(L)$, is defined to be the number of chain edges of E that overlap the lower boundary of L . Here, and throughout this section, we say that two curve segments overlap each other if they have an intersection of positive arc-length. The following is our main lemma to the proof of Theorem 5.4.

Lemma 7.1. *Given Γ and \mathcal{C} as in the beginning of this section, there exist a fully wedged set of chain edges $F^* \subseteq E(\mathcal{C})$ and an upper-disjoint set of lenses $\mathcal{L}^* \subset \mathcal{L}(\Gamma)$ such that*

$$\sum_{L \in \mathcal{L}^*} \text{weight}_{F^*}(L) \geq n^2. \tag{3}$$

The subsections that follow are devoted to the proof of this lemma. The following simple lemma directly contradicts the above lemma and completes our proof by contradiction.

Lemma 7.2. *Let $F \subseteq E(\mathcal{C})$ be any fully wedged set of chain edges and let $\mathcal{L} \subseteq \mathcal{L}(\Gamma)$ be any upper-disjoint set of lenses. Then $\sum_{L \in \mathcal{L}} \text{weight}_F(L) < n^2$.*

Proof. For each parabola P of Γ , let \mathcal{L}^P denote the set of lenses of \mathcal{L} whose upper boundary is defined by P . Since \mathcal{L} is upper-disjoint, the upper boundaries of the lenses of \mathcal{L}^P are mutually disjoint. To prove the lemma, it suffices to show that $\sum_{L \in \mathcal{L}^P} \text{weight}_F(L) < n$ for each parabola P of Γ .

Fix P . For each parabola Q such that $\text{lens}(Q, P) \in \mathcal{L}^P$, let F_Q denote the set of chain edges in F that overlap the lower boundary of $\text{lens}(Q, P)$. Thus, $|F_Q|$ is the weight of $\text{lens}(Q, P)$ with respect to F . Consider a parabola Q such that $\text{lens}(Q, P) \in \mathcal{L}^P$ and let the chain edges of F_Q enumerated as e_1, e_2, \dots, e_j in the left-to-right order, where $j = |F_Q|$. Since F is fully wedged, we have $j - 1$ distinct parabolas P_1, \dots, P_{j-1} such that each $P_i, 1 \leq i \leq j - 1$, is steeper than Q and forms a lens with Q completely contained in $\text{slab}(e_i, e_{i+1})$. These $j - 1$ parabolas form lenses with P with their upper boundaries contained in the upper boundary of $\text{lens}(Q, P)$. See Fig. 7.

Thus, for each Q such that $\text{lens}(Q, P) \in \mathcal{L}^P$, we have a set \mathcal{P}_Q of $|F_Q| - 1$ distinct parabolas. These sets \mathcal{P}_Q , for Q such that $\text{lens}(Q, P) \in \mathcal{L}^P$, are mutually disjoint, since the intersections of each parabola in \mathcal{P}_Q with P is confined in the upper boundary of $\text{lens}(Q, P)$. Collecting \mathcal{P}_Q for all Q with $\text{lens}(Q, P) \in \mathcal{L}^P$, we have altogether $(\sum_{\text{lens}(Q,P) \in \mathcal{L}^P} |F_Q|) - |\mathcal{L}^P|$ distinct parabolas. These parabolas must also be distinct from P and from each Q such that $\text{lens}(Q, P) \in \mathcal{L}^P$. Therefore, we have $\sum_{\text{lens}(Q,P) \in \mathcal{L}^P} |F_Q| - |\mathcal{L}^P| \leq n - 1 - |\mathcal{L}^P|$, or $\sum_{L \in \mathcal{L}^P} \text{weight}_F(L) \leq n - 1$. \square

7.1. Wedges and Scopes

In this subsection we introduce the notions of *wedges* and *scopes* that will form a basis for defining F^* and \mathcal{L}^* of Lemma 7.1.




Fig. 7. Proof of Lemma 7.2: parabola P and lenses \mathcal{L}^P .

We start with a basic observation on pseudoparabolas. The following lemma is obvious if our arrangement Γ consists of axis-parallel geometric parabolas: we prove it based only on the property of pseudoparabolas.

Lemma 7.3. *Let C be a cc-chain consisting of chain edges e_1, e_2, \dots, e_m in this order from left to right. Suppose that $\text{parab}(e_j)$ is the steepest among all the parabolas $\text{parab}(e_i), 1 \leq i \leq m$. Then C is strictly below $\text{parab}(e_j)$ except on e_j .*

Remark. Parabolas $\text{parab}(e_i), 1 \leq i \leq m$, in the lemma above may not be mutually distinct. The lemma implies *a posteriori*, however, that the steepest parabola $\text{parab}(e_j)$ is distinct from $\text{parab}(e_i)$ for every $i \neq j$.

Proof. By induction on m . For $m = 1$ the claim is obvious. Suppose $m \geq 2$ and let e_j be as in the lemma. Since the definition of a cc-chain implies that C is below $\text{parab}(e_j)$ in the neighborhood of e_j , it suffices to show that C does not intersect $\text{parab}(e_j)$ except on e_j . Suppose to the contrary that C intersects $\text{parab}(e_j)$ in, say, $\text{slab}(e_j, \infty)$ at vertex v . We assume that v is the leftmost such vertex if there are more than one. Thus, C is strictly below $\text{parab}(e_j)$ in $\text{slab}(e_j, v)$. Let $e_h, j < h \leq m$, be the chain edge of C that contains v . Let $\text{parab}(e_h)$ be the steepest parabola among $\text{parab}(e_{j+1}), \dots, \text{parab}(e_h)$. By the induction hypothesis, $\text{parab}(e_h)$ stays strictly above C in $\text{slab}[e_j, v]$ except on e_l . In particular, $\text{parab}(e_h)$ is strictly above $\text{right}(e_j)$ and v , which are on $\text{parab}(e_j)$. On the other hand, $\text{parab}(e_h)$ is strictly below $\text{parab}(e_j)$ on e_l . See Fig. 8. Therefore, $\text{parab}(e_h)$ must be steeper than $\text{parab}(e_j)$ contradicting the assumption that $\text{parab}(e_j)$ is the steepest among all the parabolas $\text{parab}(e_i), 1 \leq i \leq m$, and completing the induction step. \square

This observation will be used often in the following form. See Fig. 9.

Corollary 7.4. *Let P be a parabola and let u, v be points on P , with u lying to the left of v . Let C be a bounded cc-chain from u to v such that C is strictly below P in*

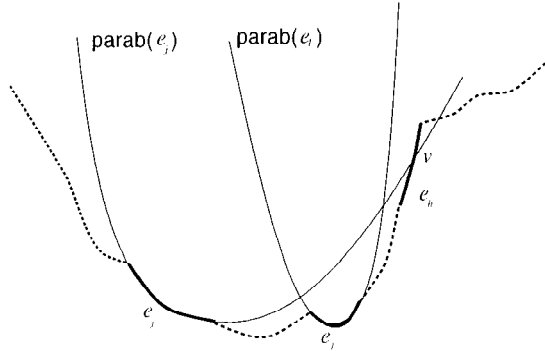


Fig. 8. Lemma 7.3: $\text{parab}(e_j)$ is supposedly the steepest.

$\text{slab}(u, v)$. Then the steepest parabola Q among those contributing chain edges to C forms a lens with P completely contained in $\text{slab}[u, v]$.

Let e be a chain edge of C . We say that a lens L of the arrangement Γ is a *left wedge* of e , if

- the upper boundary of L is defined by $\text{parab}(e)$,
- L lies to the left of e , more precisely, L is contained in $\text{slab}(-\infty, e]$, and
- the lower boundary of L overlaps some chain edge of $\text{chain}(e)$.

See Fig. 10. When L is a left wedge of e , we define the *left scope of e delimited by L* , denoted by $\text{lscope}(e, L)$, to be the open segment of $\text{parab}(e)$ in $\text{slab}(\text{left}(L), e)$. We call a left wedge L of e the *closest left wedge* of e , and denote it by $\text{lwedge}(e)$, if $\text{lscope}(e, L)$ is the minimal among all the left wedges of e . Then we define the *left scope of e* , denoted by $\text{lscope}(e)$, to be $\text{lscope}(e, \text{lwedge}(e))$, the left scope of e delimited by the closest left wedge of e ; if there is no left wedge of e , then $\text{lwedge}(e)$ is undefined and we set $\text{lscope}(e)$ to be the unbounded segment of $\text{parab}(e)$ in $\text{slab}(-\infty, e)$. The *left hinge* of e , denoted by $\text{lhinge}(e)$ is the rightmost chain edge of $\text{chain}(e)$ that overlaps the lower boundary of lens $\text{lwedge}(e)$. We define the notions of right wedges, right scopes, and right hinges analogously, with “left” and “right” swapped in the above definitions,




Fig. 9. Corollary 7.4.




Fig. 10. L is a left wedge of e and the dashed segment of $\text{parab}(e)$ is $\text{lscope}(e, L)$.

and use analogous notation: $\text{rwedge}()$, $\text{rscope}()$, and $\text{rhinge}()$. Finally, we say that a set of chain edges $E \subseteq E(C)$ is *left-sparse* (resp. *right-sparse*) if the left scope (resp. right scope) of any $e \in E$ does not overlap any $f \in E$.

To motivate the technical developments on wedges and scopes in the subsequent subsections, we sketch how these notions will be used in defining F^* and \mathcal{L}^* of Lemma 7.1. The definition of \mathcal{L}^* is via some left- or right-sparse set $E^* \subset E(C)$. Suppose E^* is left-sparse. Note that each $e \in E^*$ has a left wedge unless it is the *leftmost* chain edge of E^* on a parabola, i.e., as long as there is some $f \in E^*$ such that (f, e) is an adjacent pair of E^* , because then $\text{lscope}(e)$ must exclude f and hence is bounded. We set $\mathcal{L}^* = \{\text{lwedge}(e) \mid e \in E^* \text{ and } \text{lwedge}(e) \text{ is defined}\}$. Since the upper boundary of each lens in \mathcal{L}^* is contained in one of the mutually disjoint left scopes of the chain edges of E^* , \mathcal{L}^* thus defined is upper-disjoint as required by Lemma 7.1. \mathcal{L}^* is defined similarly when E^* is right-sparse.

The set F^* is chosen to be yet another left- or right-sparse subset of $E(C)$. Then F^* is fully wedged as required by Lemma 7.1. To see this, let (e, f) be an arbitrary adjacent pair from F^* . If F^* is left-sparse, then lens $\text{lwedge}(f)$ is contained in slab (e, f) and therefore the parabola of its lower boundary gives the parabola required in the definition of a fully wedged set. If F^* is right-sparse, then lens $\text{rwedge}(e)$ gives the required parabola.

More specific choices of E^* and F^* should be made in such a way that the weight condition (3) in Lemma 7.1 is satisfied. This is the goal of the rather involved developments in the remaining subsections.

7.2. Finding a Large Left- or Right-Sparse Set

The goal of this subsection is to find a left- or right-sparse subset of $E(C)$ that is large enough for our purposes. Recall that we are supposing $|E(C)| > 7n^{2-\epsilon}$.

Lemma 7.5. *There is a set of chain edges $E \subseteq E(C)$ with $|E| \geq 8n^{2-2\epsilon}$ that is either left- or right-sparse.*

To prove this lemma, consider the following sets of *ordered* pairs from $E(\mathcal{C})$:

$$G_{\text{left}} = \{(e, f) \in E(\mathcal{C}) \times E(\mathcal{C}) \mid \text{parab}(e) = \text{parab}(f) \text{ and } \text{right}(e) \in \text{lscope}(f)\};$$

$$G_{\text{right}} = \{(e, f) \in E(\mathcal{C}) \times E(\mathcal{C}) \mid \text{parab}(e) = \text{parab}(f) \text{ and } \text{left}(f) \in \text{rscope}(e)\}.$$

If we regard G_{left} as an *undirected* graph on node set $E(\mathcal{C})$ (ignoring the order in each pair), then $E \subseteq E(\mathcal{C})$ is left-sparse if and only if E is an independent set of G_{left} . A similar claim holds for independent sets of G_{right} . Thus, to prove Lemma 7.5, it suffices to show that the size of the maximum independent set of either G_{left} or G_{right} is at least $8n^{2-2\epsilon}$. We achieve this via showing that $G_{\text{left}} \cap G_{\text{right}}$ has a large independent set. To this end, we want to bound the number of pairs in $G_{\text{left}} \cap G_{\text{right}}$. We need some facts on wedges and scopes.

Lemma 7.6. *Let $e \in E(\mathcal{C})$ be a chain edge. Then $\text{chain}(e)$, the cc-chain of \mathcal{C} containing e , does not intersect the left or right scopes of e (and hence stays strictly below $\text{lscope}(e)$ and $\text{rscope}(e)$).*

Proof. Suppose there is an intersection of $\text{chain}(e)$ and $\text{parab}(e)$ to the left of e , i.e., in $\text{slab}(-\infty, e)$. Let v be the rightmost such intersection. We show that v is not in the left scope of e . Let C be the subchain of $\text{chain}(e)$ between v and $\text{left}(e)$, $C = \text{chain}(e) \cap \text{slab}[v, e]$. By the choice of v , C stays strictly below $\text{parab}(e)$ in $\text{slab}(v, e)$. By Corollary 7.4, there is a parabola P of Γ , contributing a chain edge to C , that forms a lens L with $\text{parab}(e)$ within $\text{slab}[v, e]$. Thus, L is a left wedge of e such that v is outside of $\text{lscope}(e, L) \supseteq \text{lscope}(e)$. An intersection of $\text{chain}(e)$ and $\text{parab}(e)$ to the right of e is dealt with similarly. □

Lemma 7.7. *Let (e, f) be a pair in $G_{\text{left}} \cap G_{\text{right}}$. Then the cc-chains $\text{chain}(e)$ and $\text{chain}(f)$ are distinct and intersect within $\text{slab}(e, f)$.*

Proof. Let $(e, f) \in G_{\text{left}} \cap G_{\text{right}}$. By the definition of G_{left} , $\text{right}(e)$ is in the left scope of f and hence, by Lemma 7.6, $\text{chain}(f)$ stays strictly below $\text{parab}(f)$ in $\text{slab}(e, f)$. Similarly, $\text{chain}(e)$ stays strictly below $\text{parab}(e) = \text{parab}(f)$ in $\text{slab}(e, f]$. In particular, $\text{chain}(e)$ is strictly above $\text{chain}(f)$ at $\text{right}(e)$ and is strictly below $\text{chain}(f)$ at $\text{left}(f)$. It follows that these two chains are distinct and must intersect in $\text{slab}(e, f)$. □

The proof above implies that, for each pair $(e, f) \in G_{\text{left}} \cap G_{\text{right}}$, there is a vertex v in $\text{slab}(e, f)$ at which $\text{chain}(e)$ crosses $\text{chain}(f)$ downward, i.e., $\text{chain}(e)$ is above $\text{chain}(f)$ on the left neighborhood of v and is below $\text{chain}(f)$ on the right neighborhood of v . We call such v a *charge vertex* of the pair (e, f) . A pair in $G_{\text{left}} \cap G_{\text{right}}$ may have more than one charge vertex. On the other hand, we have:

Lemma 7.8. *Each vertex v of Γ is a charge vertex of at most one pair in $G_{\text{left}} \cap G_{\text{right}}$.*

Proof. Suppose to the contrary that a vertex v is a charge vertex of two distinct pairs (e_1, f_1) and (e_2, f_2) in $G_{\text{left}} \cap G_{\text{right}}$. Let C_e and C_f be the cc-chains crossing at v , with




Fig. 11. Case 1 of Lemma 7.8: C_e intersects the right scope of e_1 .

C_e above C_f on the left of v . Then from the definition of a charge vertex above, $C_e = \text{chain}(e_1) = \text{chain}(e_2)$ and $C_f = \text{chain}(f_1) = \text{chain}(f_2)$. Let P_1 denote $\text{parab}(e_1) = \text{parab}(f_1)$ and P_2 denote $\text{parab}(e_2) = \text{parab}(f_2)$.

Since (e_1, f_1) and (e_2, f_2) are distinct, either $e_1 \neq e_2$ or $f_1 \neq f_2$. We assume without loss of generality that $e_1 \neq e_2$ and furthermore that e_1 lies to the left of e_2 on chain C_e .

Case 1: $P_1 = P_2$. Since e_1 and e_2 are to the left of v and f_1 is to the right of v , e_2 lies between e_1 and f_1 on P_1 . Since $(e_1, f_1) \in G_{\text{right}}$, the right scope of e_1 contains $\text{left}(f_1)$ and hence e_2 . In other words, chain C_e intersects the right scope of e_1 , contradicting Lemma 7.6. See Fig. 11.

In the remaining two cases, we assume that $P_1 \neq P_2$.

Case 2: f_1 is to the left of f_2 on chain C_f . Because $(e_1, f_1) \in G_{\text{right}}$, chain C_e stays strictly below P_1 in $\text{slab}(e_1, f_1)$. Similarly, because $(e_2, f_2) \in G_{\text{left}}$, chain C_f stays strictly below P_2 in $\text{slab}(e_2, f_2)$. Therefore, e_2 is strictly below P_1 and f_1 is strictly below P_2 , implying that P_1 intersects P_2 in $\text{slab}(e_2, f_1)$. Let u denote this intersection. Let C denote the cc-chain from $\text{right}(e_1)$ to u that first follows C_e up to e_2 and then continue on P_2 up to u . By Corollary 7.4, there must be a parabola P containing an edge of C that forms a lens with P_1 within $\text{slab}(e_1, u)$. $\text{lens}(P, P_1)$ is a right wedge of e_1 and hence $\text{left}(f_1)$ is outside of the right scope of e_1 , a contradiction to the assumption $(e_1, f_1) \in G_{\text{right}}$. See Fig. 12.

Case 3: f_1 is to the right of f_2 on chain C_f . Since (e_1, f_1) is in $G_{\text{left}} \cap G_{\text{right}}$, $\text{right}(e_1)$ is in the left scope of f_1 and $\text{left}(f_1)$ is in the right scope of e_1 . Therefore, by Lemma 7.6, both C_e and C_f stay strictly below P_1 in $\text{slab}(e_1, f_1)$. Consider a cc-chain C from $\text{right}(e_1)$ to $\text{left}(f_1)$ obtained by first following C_e up to e_2 , then follow P_2 up to f_2 , and finally follow C_f up to $\text{right}(f_1)$. See Fig. 13. If P_2 intersects P_1 within $\text{slab}(e_2, f_2)$, then, similarly, to Case 2, we have a right wedge of e_1 in $\text{slab}(e_1, f_2)$ contradicting the assumption that $\text{left}(f_1)$ is in the right scope of e_1 . Therefore, P_2 does not intersect P_1 within $\text{slab}(e_2, f_2)$ and hence chain C stays strictly below P_1 . By Corollary 7.4, there is a parabola P containing a chain edge e_P of C that forms a lens with P_1 within $\text{slab}(e_1, f_1)$.




Fig. 12. Case 2 of Lemma 7.8: $\text{left}(f_1)$ is outside of $\text{rscope}(e_1)$ because of P .

If either $e_P \in C_e$ or $P = P_2$, then $\text{lens}(P, P_1)$ is a right wedge for e_1 , contradicting the assumption that $(e_1, f_1) \in G_{\text{right}}$. Otherwise, $e_P \in C_f$ and hence $\text{lens}(P, P_1)$ is a left wedge of f_1 , contradicting the assumption that $(e_1, f_1) \in G_{\text{left}}$. \square

Since the number of vertices in Γ is at most $2 \binom{n}{2} = n(n - 1)$, Lemmas 7.7 and 7.8 imply that $|G_{\text{left}} \cap G_{\text{right}}| \leq n(n - 1)$.

We are now ready to prove Lemma 7.5. Let $\hat{E} \subseteq E(C)$ denote a maximum independent set of $G_{\text{left}} \cap G_{\text{right}}$ regarded as an undirected graph on $E(C)$. To bound the size of \hat{E} from below, we use the following well-known fact that follows from Turán’s theorem [17], [12].




Fig. 13. Case 3 of Lemma 7.8: $\text{lens}(P, P_1)$ is a left wedge of f_1 and hence $\text{right}(e_1)$ is outside of $\text{lscope}(f_1)$.

Lemma 7.9. *The size of the maximum independent set of a graph with N nodes and M edges is at least $N^2/(2M + N)$.*

Applying this bound to $G_{\text{left}} \cap G_{\text{right}}$, with $N = |E(\mathcal{C})| > 7n^{2-\varepsilon}$ (our supposition in this section), a trivial bound $N \leq n^2$, and the bound $M < n^2$ obtained above, we have

$$|\hat{E}| \geq 49n^{2(2-\varepsilon)}/(2n^2 + n^2) \geq 16n^{2-2\varepsilon}.$$

We still need to translate this bound to a bound on the size of the maximum independent set of either G_{left} or G_{right} .

Let (e, f) be an adjacent pair of \hat{E} . We say that f is *left-active* if $(e, f) \in G_{\text{left}}$ and that e is *right-active* if $(e, f) \in G_{\text{right}}$. Let E_{left} (resp. E_{right}) be obtained from \hat{E} by removing all the left-active (resp. right-active) chain edges. We claim that E_{left} (resp. E_{right}) is an independent set of G_{left} (resp. G_{right}). To see this, suppose to the contrary that $(e, f) \in G_{\text{left}}$, for some $e, f \in E_{\text{left}}$. Let e' be the chain edge of \hat{E} such that (e', f) is an adjacent pair of \hat{E} . Since the left scope of f overlaps e , it overlaps e' and hence $(e', f) \in G_{\text{left}}$. This means that f is left-active, a contradiction to the definition of E_{left} . The proof that E_{right} is an independent set of G_{right} is similar.

We next show that either $|E_{\text{left}}| \geq |\hat{E}|/2$ or $|E_{\text{right}}| \geq |\hat{E}|/2$. For each adjacent pair (e, f) of \hat{E} , we cannot have both e left-active and f right-active, because then (e, f) is in $G_{\text{left}} \cap G_{\text{right}}$ contradicting the choice of \hat{E} as an independent set of $G_{\text{left}} \cap G_{\text{right}}$. Because there are fewer than $|\hat{E}|$ adjacent pairs of \hat{E} , it follows that either fewer than $|\hat{E}|/2$ chain edges of \hat{E} are left-active or fewer than $|\hat{E}|/2$ chain edges of \hat{E} are right-active. Thus, either $|E_{\text{left}}| \geq |\hat{E}|/2$ or $|E_{\text{right}}| \geq |\hat{E}|/2$: at least one of them has size at least $8n^{2-2\varepsilon}$.

This completes the proof of Lemma 7.5. From now on, we assume without loss of generality that $|E_{\text{left}}| \geq 8n^{2-2\varepsilon}$.

7.3. Defining \mathcal{L}^*

For a set of chain edges $E \subseteq E(\mathcal{C})$, let $\text{lwedge}(E)$ denote the set of lenses $\{\text{lwedge}(e) \mid e \in E \text{ has a left wedge}\}$. Then $\text{lwedge}(E_{\text{left}})$, where E_{left} is the left-sparse set of chain edges obtained in the previous subsection, is a good candidate for \mathcal{L}^* : it is upper disjoint, its size is large, and each of its lenses contains a chain edge on its lower boundary. The latter two properties are favorable when we define F^* so as to establish the weight condition of Lemma 7.1. However, our proof requires lenses of \mathcal{L}^* to have certain additional properties: we define \mathcal{L}^* to be a subset of $\text{lwedge}(E_{\text{left}})$ as follows.

For each segment of a parabola in Γ such that its endpoints are vertices of Γ , we define the *length* of the segment to be the number of arrangement edges on the segment. We call a chain edge *short-scoped* if the length of its left scope is at most $n^{2\varepsilon}/2$. Define E_{short} to be the subset of E_{left} consisting of the short-scoped chain edges in E_{left} . Since the left scopes of the chain edges of E_{left} are mutually disjoint and there are at most n^2 arrangement edges in Γ , there are at most $2n^{2-2\varepsilon}$ chain edges of E_{left} that are not short-scoped. Therefore, we have $|E_{\text{short}}| \geq |E_{\text{left}}| - 2n^{2-2\varepsilon} \geq 6n^{2-2\varepsilon}$.

Set $t = c_1 n^{1-6\varepsilon}$, where c_1 is a sufficiently small constant, and call a lens *t-long* (resp.

t -short) if the length of the lower boundary of L is greater than t (resp. at most t). We define \mathcal{L}^* to be the subset of $\text{lwedge}(E_{\text{short}})$ consisting of t -long lenses.

Lemma 7.10. $|\mathcal{L}^*| \geq 4n^{2-2\varepsilon}$.

Proof. We first note that $|\text{lwedge}(E_{\text{short}})| \geq |E_{\text{short}}| - n \geq 5n^{2-2\varepsilon}$, because each parabola of Γ contains at most one chain edge of E_{short} that does not have a left wedge. Let \mathcal{L}' denote the subset of $\text{lwedge}(E_{\text{left}})$ consisting of t -short lenses. We claim that \mathcal{L}' forms a simple t -matching when interpreted as a set of hyperedges in the hypergraph $H(\Gamma)$. To see this, let e be an arbitrary edge of the arrangement Γ . There is at most one lens of \mathcal{L}' such that its upper boundary contains e , because \mathcal{L}' is upper-disjoint. Moreover, there are at most $t - 1$ lenses of \mathcal{L}' whose lower boundary contains e , because each of these lenses is identified by its left endvertex that must be among the $t - 1$ vertices on $\text{parab}(e)$ lining up to the left of e . Thus every arrangement edge e is contained in at most t lenses of \mathcal{L}' . Using the bound on the size of a simple k -matching in Lemma 4.5, we have $|\mathcal{L}'| \leq O(n^{5/3}t^{1/3}) = O(n^{2-2\varepsilon})$. Since the hidden constant in this O notation is increasing in c_1 , we may choose c_1 small enough so that $|\mathcal{L}'| \leq n^{2-2\varepsilon}$. It follows that $|\mathcal{L}^*| \geq |\text{lwedge}(E_{\text{short}})| - |\mathcal{L}'| \geq 4n^{2-2\varepsilon}$. \square

We define E^* to be the subset of E_{short} corresponding to \mathcal{L}^* . More formally, $E^* = \{e \in E_{\text{short}} \mid e \text{ has a } t\text{-long left wedge}\}$, so that $\mathcal{L}^* = \{\text{lwedge}(e) \mid e \in E^*\}$. Note that chain edges of E^* and lenses of \mathcal{L}^* are in one-to-one correspondence.

7.4. Defining F^*

In this subsection we define the fully wedged set of chain edges F^* in Lemma 7.1. To establish the weight condition (3), we want F^* to be sufficiently large and its elements to have many overlappings with the lower boundaries of the lenses of \mathcal{L}^* . We start with finding a large set $\tilde{F} \subseteq E(\mathcal{C})$ such that each $e \in \tilde{F}$ overlaps the boundary of some lens in \mathcal{L}^* . We later extract a fully wedged subset from \tilde{F} and define it to be F^* .

Recall the definition of left hinges. If chain edge e has a left wedge, then the left hinge of e , $\text{lhinge}(e)$, is the rightmost chain edge of $\text{chain}(e)$ that overlaps the lower boundary of the closest left wedge $\text{lwedge}(e)$ of e . Define \tilde{F} to be $\{\text{lhinge}(e) \mid e \in E^*\}$. By definition, each $e \in \tilde{F}$ overlaps the lower boundary of some lens in \mathcal{L}^* . Although we have a lower bound on $|E^*|$, this does not readily translates to a lower bound on $|\tilde{F}|$ because a chain edge may be a left hinge of more than one chain edges. For each subset E of E^* , define the *multiplicity* of E to be the maximum number of chain edges of E of which the left hinges are identical. We seek a large subset of E^* with low multiplicity, for which we need some tools.

Let e and f be distinct chain edges on a single chain C , with e lying to the left of f . We say that the pair (e, f) is *safe* if $\text{parab}(e)$ and $\text{parab}(f)$ are distinct and intersect within $\text{slab}[e, f]$.

Lemma 7.11. *Let e_1, \dots, e_m be some consecutive chain edges in a chain C , listed in the left-to-right order. If pair (e_1, e_m) is not safe, then there is some $i, j, 1 \leq i, j, \leq m$,*

such that $\text{parab}(e_i)$ and $\text{parab}(e_j)$ form a lens $\text{lens}(\text{parab}(e_i), \text{parab}(e_j))$ contained in slab $[e_1, e_m]$ whose lower boundary overlaps e_j .

Proof. The proof is by induction on m . For $m = 2$ the statement trivially holds because the pair (e_1, e_m) must be safe. Now suppose $m > 3$ and that (e_1, e_m) is not safe. First consider the case $\text{parab}(e_1) = \text{parab}(e_m)$. If pair (e_1, e_{m-1}) is safe, then setting $i = 1$ and $j = m - 1$ satisfies the lemma: $\text{lens}(\text{parab}(e_{m-1}), \text{parab}(e_1))$ is formed with its right end being $\text{right}(e_{m-1})$ and its left end being the intersection guaranteed by the definition of the safety of the pair (e_1, e_{m-1}) (and therefore lying in slab $[e_1, e_{m-1}]$). On the other hand, if pair (e_1, e_{m-1}) is not safe then the result immediately follows from the induction hypothesis.

Now suppose $\text{parab}(e_1) \neq \text{parab}(e_m)$. Without loss of generality we may assume that $\text{parab}(e_1)$ stays strictly below $\text{parab}(e_m)$ within slab $[e_1, e_m]$. Then chain C must intersect $\text{parab}(e_1)$ within slab $\text{slab}(e_1, e_m)$. Let $e_{m'}$ be the chain edge that contains the rightmost such intersection, $1 < m' < m$. If $\text{parab}(e_{m'})$ forms a lens with $\text{parab}(e_1)$ within slab $[e_1, e_{m'}]$, then we are done, setting $i = 1$ and $j = m'$. Otherwise, the pair $(e_1, e_{m'})$ is not safe and our claim directly follows from the induction hypothesis. \square

In the following, we use the fact that the size of any simple k -matching in the hypergraph $H(\Gamma)$ is at most $n^{2-4\epsilon}$. This follows from the upper bound $O(n^{5/3}k^{1/3})$ of Lemma 4.5 and our assumption $k \leq c_0 n^{1-12\epsilon}$: we exercise our reserved freedom and choose c_0 small enough so that the above bound holds.

Lemma 7.12. *There is a subset E of E^* with $|E| \geq n^{2-2\epsilon}$ whose multiplicity is at most $n^{2\epsilon}$.*

Proof. Divide each cc-chain of \mathcal{C} into subchains so that each subchain contains exactly $n^{2\epsilon}$ chain edges of E^* , discarding an incomplete subchain that may result at one end of the cc-chain. Call each of these subchains a *trail*. We have at least $|E^*|/n^{2\epsilon} - k > 3n^{2-4\epsilon}$ trails in total. We call a trail *safe* if every pair of chain edges therein is safe; otherwise it is *dangerous*. By the above lemma, to each dangerous trail we can associate a lens contained in the vertical slab spanned by the trail. Since two lenses associated with two trails of a single chain are disjoint, the collection of such associated lenses form a simple k -matching. Therefore, applying the above bound on the size of a simple k -matching, at most $n^{2-4\epsilon}$ trails are dangerous and hence at least $2n^{2-4\epsilon}$ trails are safe.

Divide each safe trail into *left* and *right parts*, so that each part contain exactly $n^{2\epsilon}/2$ chain edges of E^* . Let E be the subset of E^* consisting of chain edges in the right parts of all the safe trails. The size of E satisfies the requirement of the lemma: $|E| \geq n^{2-2\epsilon}$. It remains to show that the multiplicity of E is at most $n^{2\epsilon}$. For each $e \in E$, define the *depth* of e , denoted by $\text{depth}(e)$, to be the number of chain edges of E on the subchain starting with $\text{lhinge}(e)$ and ending with e . We claim that $\text{depth}(e) \leq n^{2\epsilon}$ for every $e \in E$. Suppose to the contrary that $\text{depth}(e) > n^{2\epsilon}$ for some $e \in E$. This implies that the subchain from $\text{lhinge}(e)$ to e contains the entire left part of the safe trail to which e belongs. Let C denote the left part of this safe trail. For each chain edge $f \in C \cap E$, pair (f, e) is safe and therefore $\text{parab}(f)$ and $\text{parab}(e)$ intersects within the slab spanned by C . Moreover,

if f_1 and f_2 are distinct chain edges in $C \cap E$, pair (f_1, \tilde{f}_2) is safe and hence $\text{parab}(f_1)$ and $\text{parab}(f_2)$ are distinct by the definition of safety. Therefore, $\text{parab}(e)$ has at least $n^{2\epsilon}/2$ vertices of Γ in the vertical slab spanned by C . Since $\text{lhinge}(e)$ is to the left of this slab, these $n^{2\epsilon}/2$ vertices are contained in the left scope of e . This contradicts the choice of E^* , which must consist of short-scoped chain edges. Therefore, $\text{depth}(e) \leq n^{2\epsilon}$ for every $e \in E$. This implies that the multiplicity of E is at most $n^{2\epsilon}$, because, if the left hinges of m chain edges of E are identical, at least one of them (the one farthest from the common left hinge) must have depth m . \square

This lemma implies that the left hinges of the chain edges of E contribute at least $|E|/n^{2\epsilon}$ distinct elements to \tilde{F} . Thus:

Corollary 7.13. $|\tilde{F}| \geq n^{2-4\epsilon}$.

We are ready to define F^* of Lemma 7.1. Let G'_{left} be the subset of G_{left} in the previous subsection defined by $G'_{\text{left}} = \{(e, f) \in G_{\text{left}} \mid e, f \in \tilde{F}\}$. Define G'_{right} analogously as a subset of G_{right} . Let E'_{left} (resp. E'_{right}) be the maximum independent set of G'_{left} (resp. G'_{right}) viewed as an undirected graph. Repeating the analysis in the previous subsection that we used for lower-bounding the size of the maximum independent set of G_{left} or G_{right} , we obtain the following.

Lemma 7.14. *Either $|E'_{\text{left}}| \geq n^{2-8\epsilon}/6$ or $|E'_{\text{right}}| \geq n^{2-8\epsilon}/6$.*

If $|E'_{\text{left}}| \geq n^{2-8\epsilon}/6$, then we set $F^* = E'_{\text{left}}$; otherwise we set $F^* = E'_{\text{right}}$. Since E'_{left} is left-sparse and E'_{right} is right-sparse, F^* is fully wedged in either case.

7.5. The Weight Condition

In this section we show that \mathcal{L}^* and F^* defined in the previous subsections satisfy the weight condition (3) of Lemma 7.1, which we list below:

$$\sum_{L \in \mathcal{L}^*} \text{weight}_{F^*}(L) \geq n^2. \tag{4}$$

To facilitate the proof, we first define a subset \mathcal{L}_1^* of \mathcal{L}^* whose members are in one-to-one correspondence with the members of F^* .

Lemma 7.15. *There exists a one-to-one mapping $\lambda: F^* \rightarrow \mathcal{L}^*$ such that each $f \in F^*$ overlaps the lower boundary of lens $\lambda(f)$.*

Proof. For each $f \in F^*$, there is at least one $e \in E^*$ such that $\text{lhinge}(e) = f$, by the definition of F^* . Choose arbitrary $e_f \in E^*$ with $\text{lhinge}(e_f) = f$ for each f and set $\lambda(f) = \text{lwedge}(e_f)$, which is in \mathcal{L}^* by the definition of E^* . To show that λ is one-to-one, let f_1, f_2 be two distinct elements of F^* . If chain edges e_{f_1} and e_{f_2} are not on the same parabola, then lenses $\text{lwedge}(e_{f_1})$ and $\text{lwedge}(e_{f_2})$ are distinct and hence we are

done. Suppose both e_{f_1} and e_{f_2} are on the same parabola P and assume without loss of generality that e_{f_1} lies to the left of e_{f_2} . We cannot have $\text{lwedge}(e_{f_1}) = \text{lwedge}(e_{f_2})$ because it would put e_{f_1} in the left scope of e_{f_2} contradicting the choice of E^* that must be left-sparse. Therefore, λ is one-to-one as required. \square

Let $\mathcal{L}_1^* \subseteq \mathcal{L}^*$ be the image $\lambda(F^*)$ of F^* . We are going to show

$$\sum_{L \in \mathcal{L}_1^*} \text{weight}_{F^*}(L) \geq n^2, \tag{5}$$

which implies our weight condition (4).

We need the following technical lemma. For each arrangement edge e of Γ , let \mathcal{L}_e denote the set of lenses of \mathcal{L}_1^* whose lower boundaries contain e . We call a lens L of \mathcal{L}_1^* *well-overlapped* if $|\mathcal{L}_e| \geq W = 12n^{8\epsilon}$ for every arrangement edge e on its lower boundary.

Lemma 7.16. *At least half of the lenses of \mathcal{L}_1^* are well-overlapped, for sufficiently large n .*

The proof of this lemma is given below. We use this lemma to prove (5). Let \mathcal{L}_2^* be the subset of \mathcal{L}_1^* consisting of well-overlapped lenses. By the above lemma and the definition of \mathcal{L}_1^* , we have $|\mathcal{L}_2^*| \geq |\mathcal{L}_1^*|/2 = |F^*|/2 \geq n^{2-8\epsilon}/12$. For each $L \in \mathcal{L}_2^*$, $\lambda^{-1}(L) \in F^*$ overlaps the lower boundaries of at least W lenses of \mathcal{L}_1^* and hence contributes a weight of at least W to the summation $\sum_{L \in \mathcal{L}_1^*} \text{weight}_{F^*}(L)$. Thus, the summation must be at least $W \cdot |\mathcal{L}_2^*| \geq n^2$.

Finally, the proof of Lemma 7.16 below completes the proof of Lemma 7.1 and hence of our theorem.

Proof of Lemma 7.16. Let \mathcal{L}_P denote the set of all lenses of \mathcal{L}_1^* whose lower boundaries are formed by parabola P . We will show that, for each P such that $|\mathcal{L}_P| \geq n^{1-8\epsilon}/24$, all but a vanishing fraction of the lenses in \mathcal{L}_P are well-overlapped. Then we will be done, because the total number of lenses that are in some \mathcal{L}_Q such that $|\mathcal{L}_Q| < n^{1-8\epsilon}/24$ is at most $n^{2-8\epsilon}/24 \leq |\mathcal{L}_1^*|/4$. Fix a parabola P with $|\mathcal{L}_P| \geq n^{1-8\epsilon}/24$. For each lens $L \in \mathcal{L}_P$ that is not well-overlapped, there is an arrangement edge e in the lower boundary of L such that $|\mathcal{L}_e| < W$, by definition: we say that e *witnesses* L .

Recall that each lens in \mathcal{L}_1^* is t -long where $t = c_1 n^{1-6\epsilon}$. Let I be an arbitrary interval on P consisting of t consecutive arrangement edges. We say that I *witnesses* a lens $L \in \mathcal{L}_P$ that is not well-overlapped, if some arrangement edge $e \in I$ witnesses L . Let S be the set of lenses witnessed by I and let S_{left} (resp. S_{right}) denote the subset of S consisting of lenses whose lower boundaries contain the leftmost (resp. rightmost) arrangement edge of I . Then we have $S = S_{\text{left}} \cup S_{\text{right}}$ because the length of the lower boundary of each $L \in S$ is at least t . Let L_0 be a lens in S_{left} such that the intersection of its lower boundary with I is minimal. Then, for every arrangement edge $e \in I$ in the lower boundary of L_0 , we have $\mathcal{L}_e \supseteq S_{\text{left}}$. Since L_0 is witnessed by some $e \in I$, this implies that $|S_{\text{left}}| < W$. Similarly we have $|S_{\text{right}}| < W$ and hence $|S| < 2W$.

Now, partition P into disjoint intervals, with each interval having t arrangement edges each. Since each lens of \mathcal{L}_P that is not well-overlapped must be witnessed by some of

these intervals, and each interval can witness at most $2W$ lenses, the number of lenses not well-overlapped is at most $2W \cdot 2n/t = 48n^{14\varepsilon}/c_1$. Since $\delta < \frac{6}{11}$ and hence $\varepsilon < \frac{1}{22}$, this number is $o(n^{1-8\varepsilon})$. \square

References

1. M. Atallah, Some Dynamic Computational Geometry Problems, *Comput. Math. Appl.* **11** (1985), 1171–1181.
2. B. Bollobas, *Graph Theory*, Springer-Verlag, New York, 1979.
3. K. Clarkson and P. Shor, Applications of Random Sampling in Computational Geometry II, *Discrete Comput. Geom.* **4** (1989), 387–421.
4. H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, New York, 1986.
5. H. Edelsbrunner and E. Welzl, On the Maximum Number of Edges in Many Faces in an Arrangement, *J. Combin. Theory Ser. A* **41** (1986), 159–166.
6. D. Eppstein, Geometric Lower Bounds for Parametric Matroid Optimization, *Proc. 27th STOC*, 1995, pp. 662–671.
7. L. Guibas and M. Sharir, Combinatorics and Algorithms of Arrangements, in *New Trends in Discrete and Computational Geometry*, J. Pach, ed., Springer-Verlag, New York, 1991, pp. 9–33.
8. D. Gusfield, Bounds for the Parametric Spanning Tree Problem, *Proc. Humbolt Conf. on Graph Theory, Combinatorics and Computing*, Utilitas Mathematica, Winnipeg, Man., 1979, pp. 173–183.
9. N. Katoh and T. Ibaraki, On the Total Number of Pivots Required for Certain Parametric Combinatorial Optimization Problems, Working Paper No. 71, Institute of Economic Research, Kobe University of Commerce, 1983.
10. N. Katoh, T. Tokuyama, and K. Iwano, On Minimum and Maximum Spanning Trees of Linearly Moving Points, *Proc. 33rd FOCS*, 1992, pp. 396–405.
11. K. Kedem and M. Sharir, An Efficient Motion Planning Algorithm for Convex Polygonal Objects Moving in 2-Dimensional Space, *Discrete Comput. Geom.* **5** (1990), 43–75.
12. J. van Lint and R. Wilson, *A Course in Combinatorics*, Cambridge University Press, Cambridge, 1992.
13. L. Lovász, On the Ratio of Optimal Integral and Fractional Covers, *Discrete Math.* **13** (1975), 383–390.
14. J. Matoušek, Lower bounds on the Length of Monotone Paths in Arrangements, *Discrete Comput. Geom.* **6** (1991), 129–134.
15. J. Matoušek, On Geometric Optimization with Few Violated Constraints, *Proc. 10th ACM on Computational Geometry*, 1994, pp. 312–321.
16. M. Sharir, On k -Sets in Arrangements of Curves and Surfaces, *Discrete Comput. Geom.* **6** (1991) 593–613.
17. P. Turán, On an Extremal Problem in Graph Theory, *Mat. Fiz. Lapok.* **48** (1941) 436–451.
18. D. J. A. Welsh, *Matroid Theory*, Academic Press, London, 1976.
19. E. Welzl, More on k -Sets of Finite Sets in the Plane, *Discrete Comput. Geom.* **1** (1986), 95–100.

Received January 17, 1996, and in revised form November 7, 1996.