On the Complexity of the Union of Fat Objects in the Plane*

Alon Efrat[†] Micha Sharir[‡]

July 16, 1998

Abstract

We prove a near-linear bound on the combinatorial complexity of the union of n fat convex objects in the plane, each pair of whose boundaries cross at most a constant number of times.

1 Introduction

Let $\mathcal{C}$ be a collection of n compact convex sets in the plane, satisfying the following properties:

(i) The objects in $\mathcal{C}$ are α-*fat*, for some fixed $\alpha > 1$; that is, for each $c \in \mathcal{C}$ there exist two concentric disks $D \subseteq c \subseteq D'$ such that the ratio between the radii of D' and D is at most α.

(ii) For any pair of distinct objects $c, c' \in \mathcal{C}$, their boundaries intersect in at most s points, for some fixed constant s.

See [12] for more details concerning fat objects in the plane.

Our goal is to derive a near-linear upper bound on the *combinatorial complexity* of the union $U = \bigcup \mathcal{C}$, where we measure the complexity by the number of intersection points between the boundaries of the sets of $\mathcal{C}$ that lie on ∂U.

*Work by Micha Sharir has been supported by the Hermann Minkowski – MINERVA Center for Geometry at Tel Aviv University, by a grant from the U.S.–Israeli Binational Science Foundation, by National Science Foundation Grants CCR-94-24398 and CCR-93-11127, and by a grant from the G.I.F. — the German Israeli Foundation for Scientific Research and Development.

†School of Mathematical Sciences, Tel Aviv University. *alone@math.tau.ac.il*

‡School of Mathematical Sciences, Tel Aviv University, and Courant Institute of Mathematical Sciences, New York University. *sharir@math.tau.ac.il*

There are not too many results of this kind. If $\mathcal{C}$ is a collection of α-fat triangles,[1] then the complexity of U is $O(n \log\log n)$ (with the constant of proportionality depending on α) [9], and this bound improves to $O(n)$ if the triangles are nearly of the same size [1]. See also [13] for additional results concerning fat polygons. If $\mathcal{C}$ is a collection of n *pseudo-disks* (arbitrary simply-connected regions bounded by closed Jordan curves, each pair of whose boundaries intersect at most twice), then the complexity of U is $O(n)$ [8]. Of course, without any additional conditions, the complexity of U can be $\Omega(n^2)$, even for the case of (non-fat) triangles. Even for fat convex objects, something like condition (ii) must be assumed, or else the complexity of the union might be arbitrarily large.

The main result of this paper is

Theorem 1.1 *The combinatorial complexity of the union of a collection $\mathcal{C}$ that satisfies conditions (i)–(ii) is $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε, α and s.*

Theorem 1.1 constitutes a significant progress in the study of the union of planar objects, an area that has many algorithmic applications, such as finding the maximal depth in an arrangement of fat objects (see [5]), hidden surface removal in a collection of fat objects in 3-space [7], and point-enclosure queries in a collection of fat objects in the plane [6]. The proof of Theorem 1.1 is given in the following sections.

2 Regular and Irregular Vertices

Let $\mathcal{C}$ be a collection of n compact simply-connected sets in the plane, each bounded by a closed Jordan curve (we refer to the sets in $\mathcal{C}$ as *Jordan regions*), and let U denote their union. We assume that these regions are in general position, so that each pair of boundaries intersect in a finite number of points and properly cross at each point of intersection, and no three boundaries have a common point. (In this subsection we make no other assumption on $\mathcal{C}$.) As already mentioned, we measure the combinatorial complexity of U by the number of vertices of the *arrangement* $\mathcal{A}(\mathcal{C})$ of $\mathcal{C}$ (i.e., points of intersection between pairs of boundaries of regions in $\mathcal{C}$) that lie on its boundary. We classify the arrangement vertices into two categories:

regular vertices: these are intersections between pairs of boundaries that intersect at only two points.

irregular vertices: these are all the other boundary intersection points.

[1] For triangles, there is an equivalent definition of fatness that requires all angles to be at least some fixed constant α_0; in [9], this is called α_0-fatness.

(In the preliminary version of the paper [4], we have referred to regular and irregular vertices as *touching* and *shattering*, respectively. The terms regular and irregular are taken from [10].)

Let $R(\mathcal{C})$ (resp. $I(\mathcal{C})$) denote the number of regular (resp. irregular) vertices of U. We use the following result of Pach and Sharir [10]:

Theorem 2.1
$$R(\mathcal{C}) \leq 2I(\mathcal{C}) + 6n - 12,$$
for $n \geq 3$.

3 Caps, Inscribed Fat Polygons, and their Properties

We now return to the case where $\mathcal{C}$ is a collection of compact convex sets satisfying the conditions (i) and (ii) in the introduction. Let $c \in \mathcal{C}$. We inscribe in c a convex polygon P_c defined as follows. We choose some constant integer parameter $t > 12$, which also satisfies
$$\frac{\alpha \sin \frac{2\pi}{t}}{1 - \frac{\pi\alpha}{t} \tan \frac{\pi}{t}} < 1.$$
and define $\theta_j = 2\pi j/t$, for $j = 0, 1, \ldots, t-1$. For each j, let $w_j = w_j(c)$ denote the (unique) point on ∂c that has a tangent (that is, a supporting line) at orientation θ_j (tangents are assumed to be oriented so that c lies to their left). We also define w'_j, for $j = 1, \ldots, t-1$, to be the point on ∂c such that the length of the portion of ∂c extending counterclockwise from w_0 to w'_j is j/t times the perimeter of c. P_c is defined to be the convex polygon whose vertices are $w_0, \ldots, w_{t-1}, w'_1, \ldots, w'_{t-1}$. (Note that P_c may have fewer than $2t-1$ vertices; this will be the case, e.g., when ∂c contains nonsmooth points whose tangent orientations span a sufficiently large interval.) The difference $c \setminus P_c$ is the union of at most $2t-1$ *caps* of c, where a cap is an intersection of c with a halfplane. The *chord* of a cap is the intersection of c with the line bounding the corresponding halfplane. An illustration of such an inscribed polygon and of the corresponding caps is shown in Figure 1.

Lemma 3.1 *The polygons P_c are α'-fat, for*
$$\alpha' = \frac{\alpha}{1 - \frac{\pi\alpha}{t} \tan \frac{\pi}{t}}.$$

Proof: Since c is α-fat, there exist two concentric disks $D_1 \subseteq c \subseteq D_2$, with respective radii r_1, r_2, such that $r_2 \leq \alpha r_1$. Clearly, $P_c \subseteq D_2$. Let K be one of the caps that constitute $c \setminus P_c$, and assume that D_1 intersects the chord pq of K. It must do so at two points, or else its interior would have contained p or q, contradicting the assumption

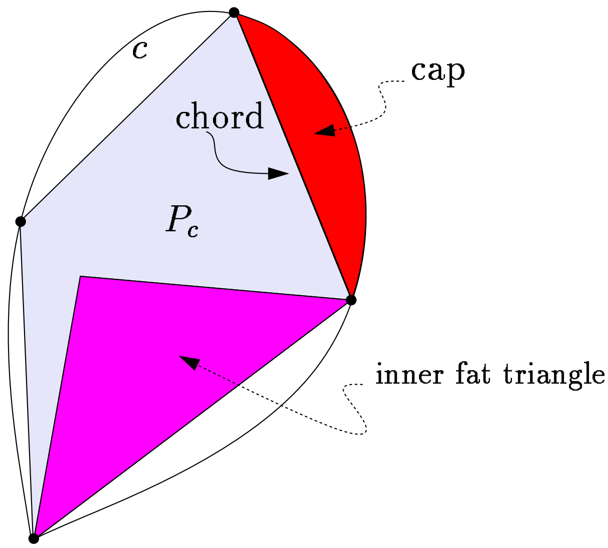

Figure 1: The inscribed polygon P_c and the corresponding caps; one inner fat triangle is also illustrated.

that $D_1 \subseteq c$. By definition, there exist two tangents to c, τ_p at p and τ_q at q, whose orientations differ by at most $2\pi/t$, and the distance pq is smallest than σ/t, where σ is the perimeter of c. Let d denote the distance from the center O of D_1 to pq. It is easy to see that $r_1 - d$ is at most the height to pq in the triangle bounded by pq, τ_p and τ_q (see Figure 2), and a simple exercise shows that this height is at most $\frac{pq}{2}\tan\frac{\pi}{t}$. Hence

$$r_1 - d \leq \frac{pq}{2}\tan\frac{\pi}{t} < \frac{\sigma}{2t}\tan\frac{\pi}{t} \leq \frac{2\pi\alpha r_1}{2t}\tan\frac{\pi}{t} = \frac{\pi\alpha r_1}{t}\tan\frac{\pi}{t},$$

where the last inequality follows from the fact that $c \subseteq D_2$. Hence

$$d \geq r_1\left(1 - \frac{\pi\alpha}{t}\tan\frac{\pi}{t}\right).$$

This implies that the disk concentric with D_1 and having radius $r_1\left(1 - \frac{\pi\alpha}{t}\tan\frac{\pi}{t}\right)$ is contained in P_c, and this completes the proof of the lemma. $\Box$

Let $c \in \mathcal{C}$, and let O denote the common center of two disks $D_1 \subseteq P_c \subseteq D_2$, such that their respective radii r_1, r_2 satisfy $r_2 \leq \alpha' r_1$. Let pq be an edge of P_c. The convexity of P_c and the fact that $D_1 \subseteq P_c$ are easily seen to imply that the angle Opq must be at least the angle β between Op and the tangent to D_1 from p, which satisfies $\sin\beta = r_1/|Op| \geq r_1/r_2 \geq 1/\alpha'$. Similarly, the angle Oqp must also be at least β. It follows that we can find a point v inside Opq, such that all the angles of the triangle vpq are at least

$$\beta_0 = \min\{\arcsin(1/\alpha'), \pi/3\}.$$

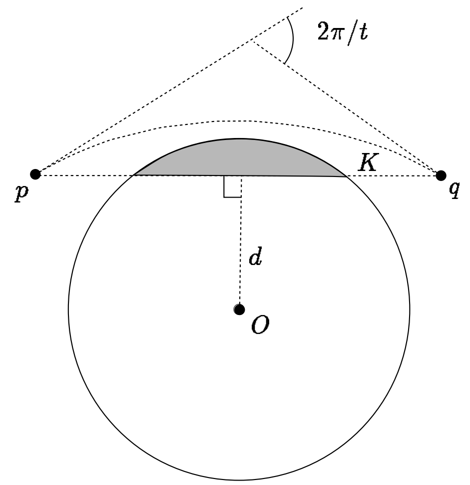

Figure 2: The proof of Lemma 3.1

Note that, by assumption, $\beta_0 > 2\pi/t$.

We repeat this analysis to each edge of each polygon, and replace the polygons P_c by the collection of resulting triangles vpq. We refer to these triangles as *inner fat triangles*. Let $\mathcal{T} = \mathcal{T}(\mathcal{C})$ denote the collection of inner fat triangles. Clearly, $|\mathcal{T}| \leq (2t-1)n$. As an immediate consequence of [9], we have:

Lemma 3.2 *The union $U_{\mathcal{T}}$ of the triangles in $\mathcal{T}$ has $O(n \log\log n)$ vertices.*

Let v be an irregular vertex of ∂U, incident to two sets $a, b \in \mathcal{C}$. Let K_a, K_b be the respective caps of a, b that contain v, and let $p_a q_a$, $p_b q_b$ denote their respective chords. Consider the convex set $R = K_a \cap K_b$.

Lemma 3.3 *At least one of the chords $p_a q_a$, $p_b q_b$ meets ∂R.*

Proof: Indeed, suppose to the contrary that both chords are disjoint from R. It follows that $R = a \cap b$, and that ∂R contains at least four points of intersection between ∂a and ∂b. Moreover, let O be an interior point of R, and consider ∂K_a and ∂K_b as graphs of two respective functions $r = K_a(\theta)$, $r = K_b(\theta)$, in polar coordinates about O. Note that ∂R is the graph of the pointwise minimum of K_a and K_b. There is an angular interval I_a over which $K_a(\theta)$ is attained at the chord of K_a, and a similar interval I_b for the chord of K_b. These intervals must be disjoint, or else ∂R would overlap one of these chords, contrary to assumption. See Figure 3.

Let u (resp. w) denote the first vertex of ∂R that we encounter as we rotate about O clockwise (resp. counterclockwise) from I_a (clearly, no vertex of ∂R has an orientation in I_a). In the angular interval that runs counterclockwise from u to w, the boundary of R is attained by ∂b. Moreover, as we traverse, in counterclockwise direction, the portion of ∂b that lies on ∂K_b, we first encounter u and then w, and the reverse order is obtained along ∂a. See Figure 3.

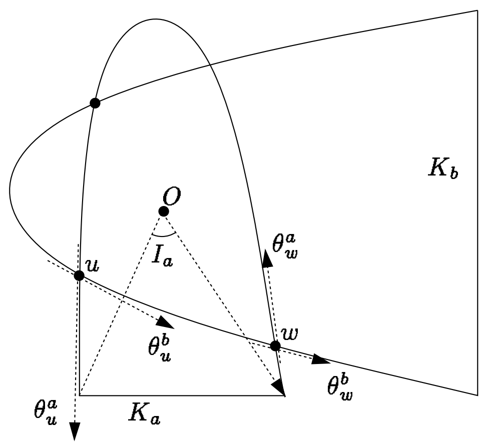

Figure 3: Two intersecting caps without a chordal intersection

Let θ^a_u, θ^a_w denote the orientations of the tangents to a at u and w, respectively, and let θ^b_u, θ^b_w denote the corresponding tangent orientations for b. (If any of these tangents is not unique, we fix an arbitrary tangent among those that are available.) The circular counterclockwise order of these four orientations is $(\theta^a_u, \theta^b_u, \theta^b_w, \theta^a_w)$, and they partition the circular range of orientations into four angular intervals that we denote by (θ^a_u, θ^b_u), (θ^b_u, θ^b_w), (θ^b_w, θ^a_w), and (θ^a_w, θ^a_u). Each of the second and fourth intervals has length at most $2\pi/t$ (since the endpoints of any of these intervals are two tangent orientations within a single cap), and each of the first and third intervals has length at most π (the total amount by which the tangent to a convex set can turn at a fixed point of its boundary is at most π). It follows that each of the lengths of the first and third intervals is at least $\pi - 4\pi/t > 2\pi/3$.

We now repeat the whole analysis in the last two paragraphs by interchanging a and b. This yields two vertices u', w' of ∂R, such that the turning angle of the tangents to R at each of these vertices is also greater than $2\pi/3$. It is easily verified that among the vertices u, w, u', w' there exist at least three distinct vertices, or else ∂a and ∂b would have intersected at only two points, contrary to assumption. We have thus obtained at least three vertices of R such that the turning angle of the tangents at each of them is greater than $2\pi/3$, which is impossible, because the overall turning angle for a convex set is 2π. This contradiction completes the proof of the lemma. □

Lemma 3.4 *Let K_a be a cap of some set $a \in \mathcal{C}$, with chord e_a, and let Δ_b be an inner fat triangle in $\mathcal{T}$, obtained from the polygon P_b, for some $b \in \mathcal{C}$, such that the chord e_b of Δ_b crosses ∂K_a. Then one of the following cases must occur:*

(i) e_a crosses $\partial \Delta_b$ (as in Figure 4(i)).

(ii) K_a *contains a vertex of* Δ_b *that is an endpoint of* e_b *(as in Figure 4(ii)).*

(iii) Δ_b *contains a vertex of* K_a *(as in Figure 4(iii)).*

(iv) ∂K_a *and* $\partial \Delta_b$ *cross exactly twice, at two points that lie on* ∂a *and on* e_b*, and* e_a *is disjoint from* $K_a \cap \Delta_b$*. Furthermore, let* K_b *denote the cap of* b *that shares the same chord* e_b *with* Δ_b*. Then either* K_b *contains an endpoint of* e_a *(as in Figure 4(iv.a)), or* ∂a *and* ∂b *intersect only twice (as in Figure 4(iv.b)).*

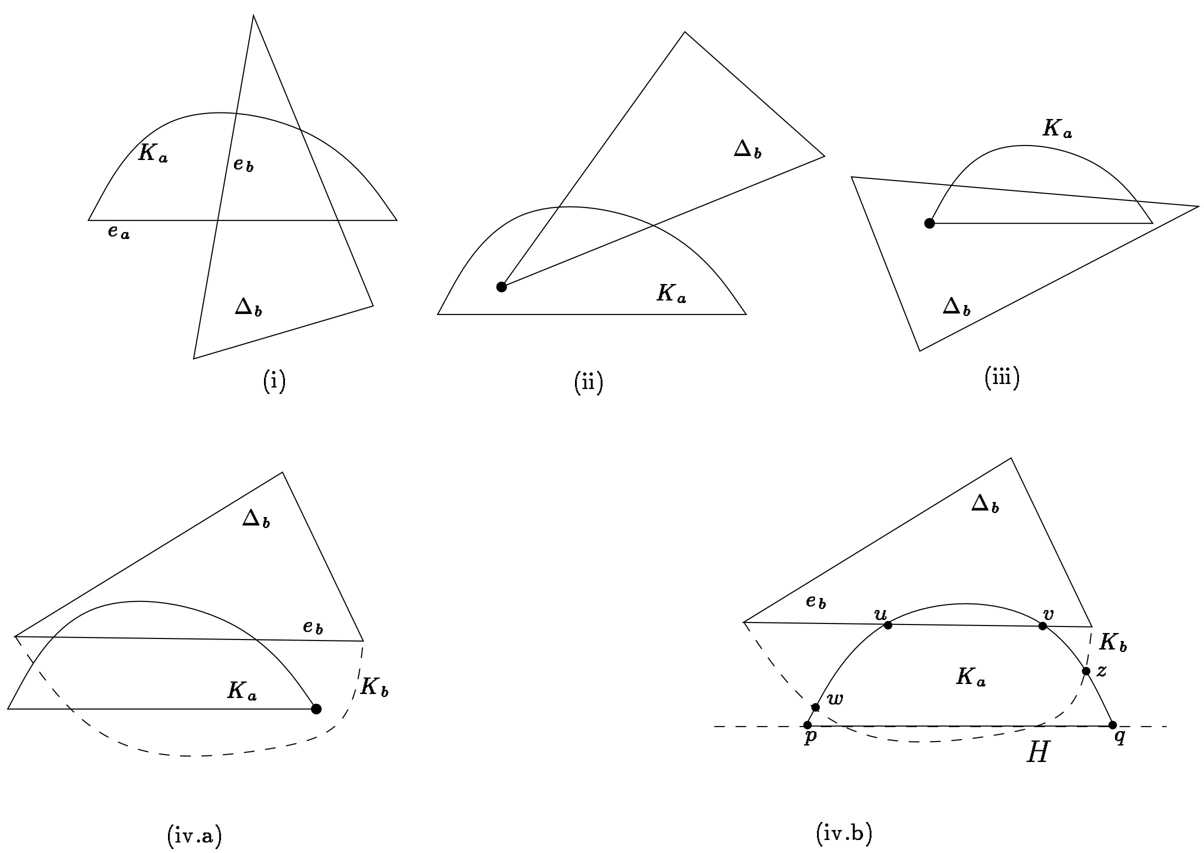

Figure 4: Illustrating the various cases in Lemma 3.4

Proof: Suppose that cases (i) and (ii) do not occur. That is, e_a does not cross $\partial \Delta_b$ and no vertex of Δ_b lies in K_a. Then e_b must intersect ∂K_a at two points, u, v, both lying on ∂a. Therefore e_b splits K_a into two subregions, the region K_a' that contains e_a, and the complementary region K_a''. Denote the range of the orientations of the tangents to a at the points of K_a by (θ_0, θ_0'), where $\theta_0 < \theta_0' \le \theta_0 + 2\pi/t$. Clearly, the orientations of e_a and of e_b also lie in this range. Two cases can arise:

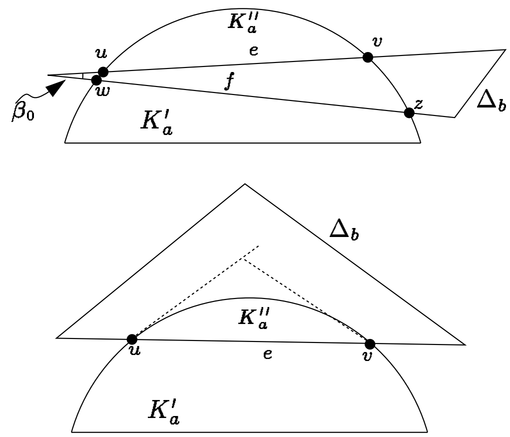

[htbp]

Figure 5: Two patterns of intersection of a cap K_a and an inner fat triangle Δ_b

(1) Δ_b overlaps K'_a and is disjoint from K''_a (see Figure 5(i)): If K'_a is fully contained in Δ_b then u and v are the only two points of intersections between ∂K_a and $\partial \Delta_b$, and, moreover, Δ_b contains both vertices of K_a, so we are in case (iii). Otherwise, since, by assumption, Δ_b does not intersect e_a and does not have a vertex inside K'_a, one of its other edges, f, must also cross ∂K_a twice, at two points w, z, lying on ∂a, so that the four points w, u, v, z appear in this order along ∂K_a. In this case the orientation of f also lies in the range (θ_0, θ'_0), and thus the angle between e and f, which is $\geq \beta_0$, is at most $2\pi/t$, a contradiction.

(2) Δ_b overlaps K''_a and is disjoint from K'_a (see Figure 5(ii)): We claim that in this case Δ_b fully contains K''_a, so u and v are the only two intersection points of ∂K_a and $\partial \Delta_b$. Since the orientations of e_b and of the tangents (or, rather, any tangents) to a at u and at v all lie in the range (θ_0, θ'_0), it follows that the angles between e and these tangents are both at most $2\pi/t$. However, the angles of Δ_b at the endpoints of e are both $\geq \beta_0$, and are therefore larger. It follows that the triangle bounded by e and by two such tangents is fully contained in Δ_b, from which the claim follows readily.

Finally, suppose that K_b does not contain any of the endpoints e_a. Let p and q be the endpoints of e_a, so that p, u, v, q appear in this order along ∂a. Then the portion of ∂K_b along ∂b must cross the portion of ∂K_a along ∂a at least twice, at one point w between p and u and at another point z between v and q (see Figure 4(iv.b)). We claim that w and z are the only two intersection points of ∂a and ∂b. Indeed, suppose, with no loss of generality, that e_a lies along the x-axis and that K_a lies above it. Then $\gamma_a \equiv \partial a \cap K_a$ is a downward-concave x-monotone arc. Moreover, the absolute value of the orientation of e_b is at most $2\pi/t$, so the orientation of any tangent to $\gamma_b \equiv \partial b \cap K_b$ has absolute value $\leq 4\pi/t$, which is easily seen to imply that γ_b is also x-monotone and downward-convex. It follows that γ_a and γ_b cross each other exactly twice (at w and z). We claim that there can be no other point of intersection between ∂a and ∂b. Indeed, any such point must lie either in the halfplane below e_a or in the halfplane above e_b. Consider the halfplane H lying below e_a (the second case is treated in a fully symmetric manner). It is easy to see that any such intersection must lie on γ_b. However, if γ_b reaches H it must cross e_a twice. Arguing as above, it follows that the portion of γ_b in H is fully contained in the inner fat triangle of P_a that has e_a as a chord, and hence it cannot intersect ∂a at all. This shows that condition (iv) holds, and thus completes the proof of the lemma. (Note that these arguments also imply that, in any configuration of case (iv), $\partial a \cap K_a$ and $\partial b \cap K_b$ can intersect in at most two points; they intersect in one or zero points if and only if K_b contains an endpoint of e_a.) $\Box$

4 The Proof of Theorem 1.1

The proof proceeds by induction on n, keeping ε, α and s fixed. Let $F(n)$ denote the maximum number of vertices of the union of any collection $\mathcal{C}$ of n compact convex objects, satisfying conditions (i)–(ii) of the introduction. We will show that $F(n) \le Bn^{1+\varepsilon}$, where B is a sufficiently large constant, depending on ε, α and s. By choosing B sufficiently large, this will hold for all $n \le n_0$, where the value of n_0 will be defined below. Suppose then that $n > n_0$ and that the claim holds for all $n' < n$.

Step I. For each $c \in \mathcal{C}$, let Q_c denote a smallest axis-parallel square enclosing c. We construct a two-dimensional hereditary segment tree on the collection $\mathcal{Q} = \{Q_c \mid c \in \mathcal{C}\}$, as follows. We construct a one-dimensional segment tree T_1 on the x-projections of the squares in $\mathcal{Q}$. We make T_1 *hereditary*, as in [2], by propagating a square Q_c that is normally stored at some node ξ of T_1 to all ancestors of ξ. In this manner, each node ξ of T_1 stores two lists: the standard list $L_1(\xi)$ of squares stored at ξ (we refer to these squares as *long*), and a list $S_1(\xi)$ of squares that were propagated to ξ from its (proper) descendants. The total length of all these lists is still $O(n \log n)$.

We now take each node ξ of T_1, and construct a secondary hereditary segment tree $T_2^{(\xi)}$ on the y-projections of the squares in $L_1(\xi) \cup S_1(\xi)$. Again, each node η of any secondary tree stores two lists: the standard list $L_2(\eta)$ of 'long' squares, and a list $S_2(\eta)$ of 'short' squares, propagated from the proper descendants of η. The total size of the structure is $O(n \log^2 n)$.

Let v be a vertex of the union, lying on the boundaries of two sets $a, b \in \mathcal{C}$. We take the leaf ζ of T_1 whose x-interval contains the x-coordinate of v, and consider the path from ζ to the root of T_1. There is a unique node ξ on that path such that one of Q_a, Q_b is stored at $L_1(\xi)$ and the other square is stored at $L_1(\xi) \cup S_1(\xi)$. Repeating this for the secondary tree $T_2^{(\xi)}$ and the y-coordinate, we obtain a unique node η of $T_2^{(\xi)}$ (in fact of the whole structure) such that v lies in the rectangle R_η, defined as the cartesian product of the x-interval associated with ξ and the y-interval associated with η, and such that one of the squares Q_a, Q_b is stored at $L_2(\eta)$ and the other is stored at $L_2(\eta) \cup S_2(\eta)$.

Our strategy is thus to iterate over all vertices η of all the secondary trees $T_2^{(\xi)}$, and, for each fixed ξ, η, prove a near-linear upper bound on the number of vertices v of U, incident to a pair of objects $a, b \in \mathcal{C}$, satisfying

- $v \in R_\eta$;
- both squares Q_a, Q_b are in $L_2(\eta) \cup S_2(\eta)$ (and thus also in $L_1(\xi) \cup S_1(\xi)$).
- At least one of these squares is in $L_1(\xi)$, and at least one is in $L_2(\eta)$.

We prune the sets $L_2(\eta)$, $S_2(\eta)$, so as to retain only squares Q_a whose object a intersects R_η. Clearly, the a and b above are not pruned by this rule. We continue to use the same notation $L_2(\eta)$, $S_2(\eta)$, for the pruned sets.

Let ξ and η be fixed.

Lemma 4.1 *(a) Suppose that the height of R_η is larger than or equal to its width. Then there exists a set P_η of $O(1)$ points, all lying in the rectangle R'_η obtained by scaling up R_η by a factor of 2 about its center, such that any $a \in \mathcal{C}$ with $Q_a \in L_2(\eta)$ has a nonempty intersection with P_η.*

(b) Suppose that the height of R_η is smaller than its width. Then there exists a set P_η of $O(1)$ points, as above, such that any $a \in \mathcal{C}$ with $Q_a \in L_1(\xi) \cap (L_2(\eta) \cup S_2(\eta))$ has a nonempty intersection with P_η.

Proof: Consider the proof of (a). Let $a \in \mathcal{C}$ be such that $Q_a \in L_2(\eta)$. Since (i) a intersects R_η, (ii) the y-projection of Q_a contains that of R_η, (iii) the y-projection of R'_η is at least as large as its x-projection, and (iv) a is α-fat, it follows that the area of $a \cap R'_\eta$ is at least some fixed portion of the area of R'_η (see [6]). Hence, if we place a sufficiently dense grid of $O(1)$ points within R'_η, at least one of them will lie in a. The proof of (b) is fully symmetric. $\Box$

Let P_η be the point set yielded by Lemma 4.1, and fix a point p in P_η. Let $\mathcal{C}_1(p)$ (resp. $\mathcal{C}_2(p)$) denote the collection of sets $a \in \mathcal{C}$ that contain (resp. do not contain) p, and whose enclosing squares Q_a are in $L_2(\eta) \cup S_2(\eta)$. The preceding analysis implies that each vertex v of U that satisfies the above conditions has at least one $p \in P_\eta$ such that one of the sets whose boundaries contain v lies in $\mathcal{C}_1(p)$ and the other set lies in $\mathcal{C}_1(p) \cup \mathcal{C}_2(p)$.

All these reductions imply that it suffices to solve the following problem: We are given two families $\mathcal{C}_1$, $\mathcal{C}_2$ of α-fat convex objects in the plane, each pair of whose boundaries intersect at most s times. We are also given that all the objects in $\mathcal{C}_1$ contain a fixed point, which from now on we take to be the origin. We want to obtain a near-linear bound for the number of vertices of the union of $\mathcal{C}_1 \cup \mathcal{C}_2$ that lie on the boundary of at least one set in $\mathcal{C}_1$.

Put $m_1 = |\mathcal{C}_1|$ and $m_2 = |\mathcal{C}_2|$. Let U_1 denote the union of $\mathcal{C}_1$. If we represent the boundary of every $c \in \mathcal{C}_1$ as the graph of a function $r = r_c(\theta)$ in polar coordinates, the boundary of U_1 is the graph of the upper envelope of these functions. Since any pair of these functions intersect in at most s points, the number of vertices of U_1 is at most $\lambda_s(m_1)$ [11].

For each $a \in \mathcal{C}_2$, we construct an inner fat inscribed polygon P_a, as in Section 3, with the additional proviso that the vertices of P_a also include the two points on ∂a that have extreme clockwise and counterclockwise orientations. (Since a does not contain the origin, these points are well defined, except that any of them may actually be replaced by a radially-directed segment on ∂a; in this case, the two endpoints of such an interval are assumed to be vertices of P_a.)

Let U_2 denote the union of the polygons P_a, for $a \in \mathcal{C}_2$. By the result of [9], the number of vertices of U_2 is $O(m_2 \log\log m_2)$.

Step II. As an intermediate step, we bound the complexity of $U^* = U_1 \cup U_2$. It suffices to bound the number of 'mixed' vertices of U^*, namely, vertices that lie on both ∂U_1 and ∂U_2.

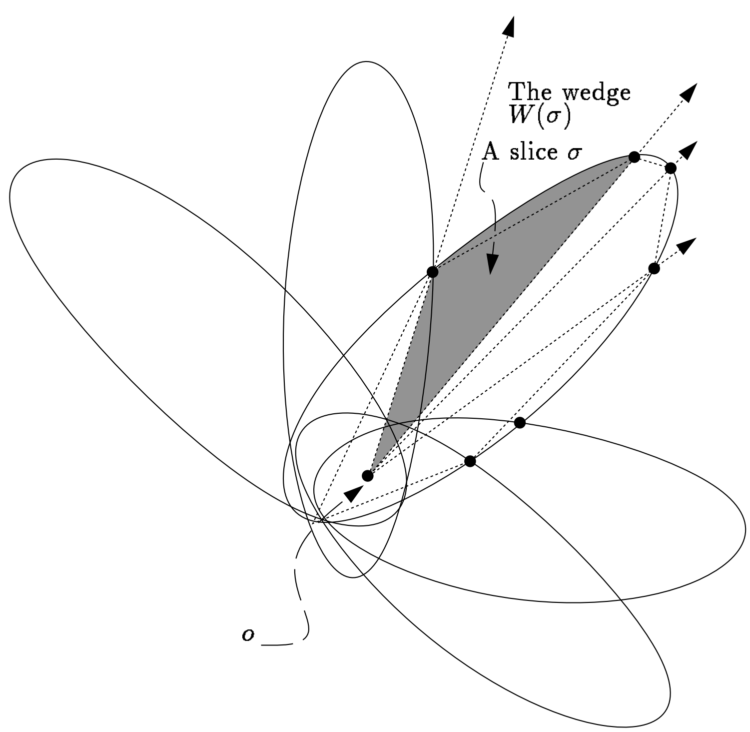

Figure 6:

To begin with, we take each $a \in \mathcal{C}_1$ and modify its inner fat inscribed polygon P_a by adding all the vertices of ∂U_1 that lie on ∂a as vertices of P_a. We obtain a collection of new polygons P^*_a, which may now have more than a constant number of vertices on each. Still, the overall number of their vertices is $O(\lambda_s(m_1))$.

Next, we take each point w which is either a vertex of U_1 or a vertex of some P_a, and connect it by a straight segment to the origin. These segments partition U_1 into 'slices', each bounded by two of these segments and by a portion of the boundary of a single set in $\mathcal{C}_1$. Each slice is further subdivided, by the *chord* connecting its two vertices, into a cap (as in Section 3) and a triangle with the origin as a vertex. We also define the *wedge* $W(\sigma)$ of a slice σ to be the positive hull of σ (it is the wedge with the origin as apex and with two bounding rays containing the segments bounding σ).

Let v be a mixed vertex, lying on the boundary of a set $a \in \mathcal{C}_1$ and on the boundary of a polygon P_b for some $b \in \mathcal{C}_2$. Let σ be the slice of U_1 that contains v, and let e be the edge of P_b that contains v. Suppose first that P_b and the origin lie on opposite sides of the line containing e. If e intersects $\sigma \cap \partial a$ at two points, then Lemma 3.4(iv.b) implies that P_b (or, more precisely, the inner fat triangle of P_b having e as an edge) and σ intersect in just two points (one of which is v), so v is a regular vertex of U^* (viewed as the union of the slices and the inscribed fat polygons of the sets in $\mathcal{C}_2$). By Theorem 2.1, the number of such vertices is at most $2I + 6(m_1 + m_2) - 12$, where I is the number of irregular vertices of U^*. It thus suffices to bound the number of irregular mixed vertices of U^*, so the above case can be ignored. If e intersects $\sigma \cap \partial a$ only at v, then it must intersect the chord of σ, or end inside the cap. This means that if we follow e from v into σ, we encounter there a vertex of the union $S_1 \cup U_2$, where S_1 is the star-shaped polygon composed of all the slice chords. We then charge v to the first such vertex that we encounter, and note that this charging is unique.

Lemma 4.2 *The complexity of $S_1 \cup U_2$ is $O((\lambda_s(m_1) + m_2)\log\log n)$.*

Proof: We take the modified inscribed fat polygons P_a^*, and decompose each of them into fat triangles, as in Section 3. It is clear that ∂S_1 is contained in the boundary of the union of these triangles. It follows that every vertex of $S_1 \cup U_2$ is also a vertex of the union of $O(\lambda_s(m_1) + m_2)$ α'-fat triangles, and the lemma thus follows from [9]. $\Box$

Hence, the number of vertices v in the preceding case is $O((\lambda_s(m_1)+m_2)\log\log n)$.

We may therefore assume that P_b and the origin lie on the same side of the line containing e. The intersection of e and of $\sigma \cap \partial a$ consists of one or two points. Suppose first that there is only one point of intersection, namely v. We trace e from v into σ, and note that the line containing e must intersect the boundary of the cap of σ at another point v'. If we encounter a vertex w of U_2 (which can be the endpoint of e or an earlier point on e) before reaching v', we charge v to w and note that w can be charged at most twice. The overall number of such vertices v is $O(m_2 \log\log m_2)$.

Otherwise we reach v', which necessarily lies on the chord of σ. As above, v' is a vertex of the union $S_1 \cup U_2$, and we can charge v to v', note that the charging is unique, and conclude, by Lemma 4.2, that the number of vertices v in this subcase is $O((\lambda_s(m_1) + m_2)\log\log n)$.

Suppose then that e intersects $\sigma \cap \partial a$ at two points, v and v', and that its portion within σ contains no vertex of U_2 (otherwise we can charge v to the first such vertex along e, as above). Hence v' is also a vertex of U^*. If e terminates within the wedge $W(\sigma)$ of σ, we can charge v and v' to such an endpoint. This charging is 'almost unique': Since these wedges are pairwise openly disjoint, the endpoint determines e (there are in fact two choices for e) and σ (there can be two choices for σ if the endpoint lies on a wedge boundary), so v and v' are also determined (at worst there can be four such pairs). The total number of vertices v in this subcase is thus $O(m_2)$. We may thus assume that e fully crosses the wedge $W(\sigma)$.

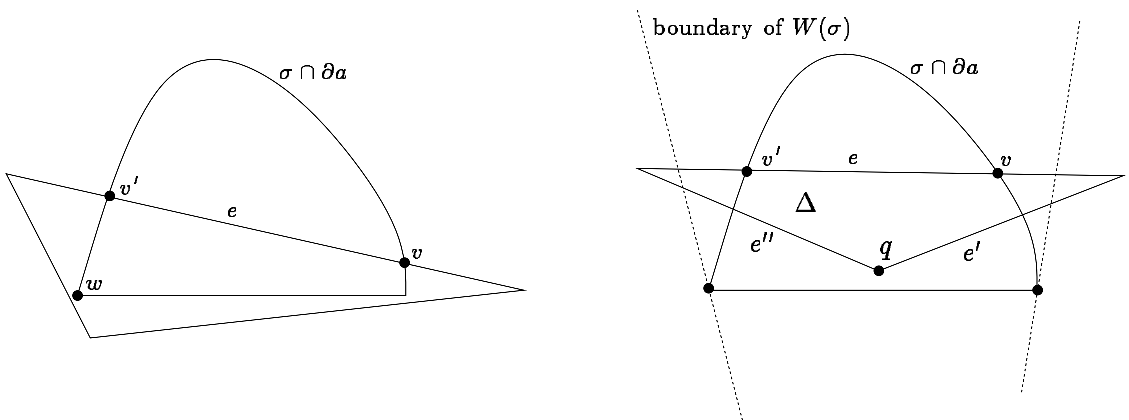

Figure 7: Two cases where P_b and the origin lie on the same side of the line containing e

Trace $\sigma \cap \partial a$ from v and from v' into P_b. If we reach along one of these arcs an endpoint w of $\sigma \cap \partial a$, we charge v and v' to w, note that w can be charged at most twice, and conclude that the number of vertices v in this subcase is $O(\lambda_s(m_1))$. Otherwise, each of these arcs is crossed by another edge of the inner fat triangle Δ bounded by e. The analysis in Lemma 3.4 implies that it is impossible that such an edge $e' \neq e$ crosses both arcs, and it is impossible for such an e' to cross the same arc twice (in both cases one of the angles of Δ would have to be at most $2\pi/t < \beta_0$). Since the common endpoint z of e and e' lies outside the wedge $W(\sigma)$, it follows that as we trace e' from z, we first meet the wedge boundary (still outside σ) and then cross $\sigma \cap \partial a$. Since this must also hold for the third edge e'' of Δ, it follows that the vertex q of Δ where e' and e'' meet must lie inside σ. We can thus charge v and v' to q, note that the charging is unique (knowing q we also know σ and Δ), and conclude that the number of vertices v in this subcase is $O(m_2)$.

We have thus shown that the complexity of U^* is $O((\lambda_s(m_1) + m_2) \log\log n)$.

Step III. We finally turn to estimate the number of mixed vertices of the union U of $\mathcal{C}_1 \cup \mathcal{C}_2$.

First, we take all the mixed vertices of U^* and connect them to the origin, thereby splitting some slices of U_1 into subslices. The total number of new slices is $O((\lambda_s(m_1)+m_2)\log\log n)$. Second, we concentrate only on irregular vertices of the union. We will later exploit Theorem 2.1 to take into account regular vertices as well.

Let v be an irregular vertex of U, lying on the boundary of a set $a \in \mathcal{C}_1$ and a set $b \in \mathcal{C}_2$. Let σ be the (new) slice of U_1 containing v. With no loss of generality, we may assume that when we follow $\partial\sigma$ from v in the counterclockwise direction, we enter into b. The complementary type of irregular vertices will be handled by a fully symmetric analysis.

We will further classify these vertices into two subcategories: For each $b \in \mathcal{C}_2$, the two tangents from the origin to b divide ∂b into two arcs, one being 'visible' from the origin and the other being 'invisible'. We refer to these portions as the *lower boundary* and *upper boundary* of b, respectively. We will consider separately vertices v as above that lie on the lower boundaries of the sets of $\mathcal{C}_2$ and vertices that lie on the upper boundaries.

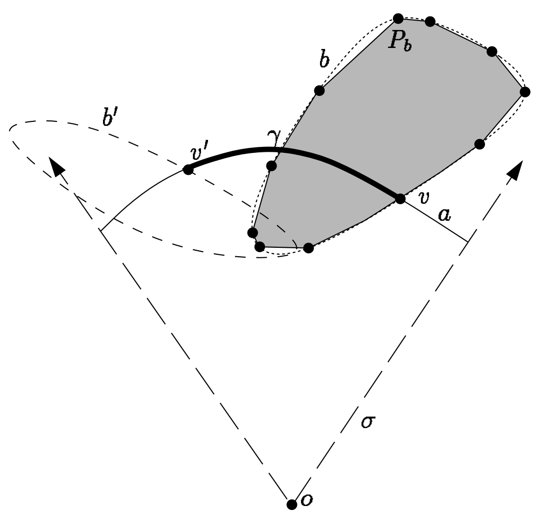

Figure 8: The case of vertices on lower boundaries

Vertices on lower boundaries. Let v and v' be two vertices as above, lying within a single new slice σ and incident to the lower boundaries of two respective and distinct sets b, $b' \in \mathcal{C}_2$, see Figure 8. With no loss of generality, assume that v lies clockwise to v'. Consider the portion γ of ∂a between v and v'. The arc γ partially overlaps the interior of b near v and it must cross ∂b again. If it only crosses the lower boundary of b then it is easily checked that ∂a and ∂b intersect only twice, so v is a regular vertex, contrary to assumption. Hence, γ also crosses the upper boundary of b. But then, by construction, γ must also cross the boundary of P_b, and thus it must contain a vertex of U^*, which implies, by construction, that v and v' cannot belong to the same new slice, again a contradiction.

We thus conclude that any new slice can have at most one vertex of the union of the type under consideration, so the number of these vertices is $O((\lambda_s(m_1)+m_2)\log\log n)$.

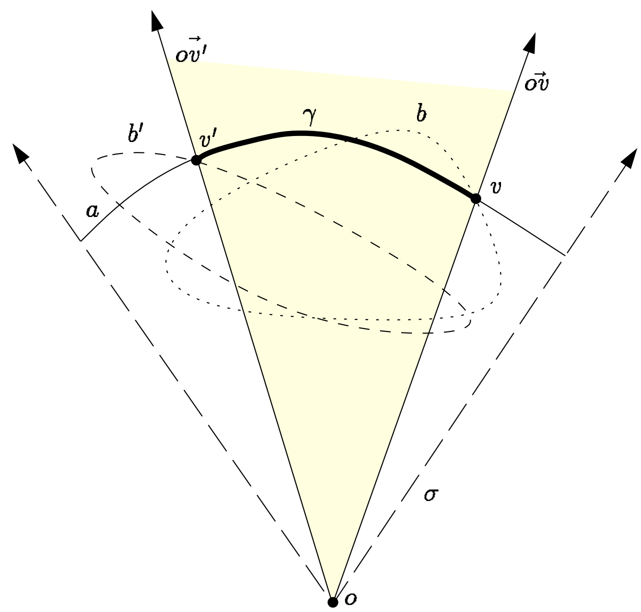

Figure 9: ∂b and $\partial b'$ must cross inside the shaded wedge

Vertices on upper boundaries. Let us fix the new slice σ, as above, and consider the number N_σ of sets in $\mathcal{C}_2$ whose upper boundaries are incident to vertices v of the above kind that lie on $\sigma \cap \partial a$. If such a set b has a supporting line that passes through the origin and is contained in the wedge $W(\sigma)$, then we charge the corresponding vertex or vertices to that line. This charging is almost unique, since the line determines both the set b and the slice σ. Hence the number of vertices v of this kind is at most sm_2. We can therefore exclude such pairs (b, σ) from our analysis.

Let v and v' be two vertices as above, lying within a single new slice σ and incident to the upper boundaries of two respective and distinct sets b, $b' \in \mathcal{C}_2$; see Figure 9. With no loss of generality, assume that v lies clockwise to v'. Consider the portion γ of ∂a between v and v'. As above, the arc γ partially overlaps the interior of b near v and it must cross ∂b again. If it crosses the lower boundary of b then, arguing as above, γ must contain a vertex of U^*, which is impossible. Hence γ only crosses the upper boundary of b. The same argument also implies that γ cannot cross the lower boundary of b'.

We claim that ∂b and $\partial b'$ must intersect within the wedge bounded by the two rays $\vec{ov}$, $\vec{ov'}$, emerging from the origin towards v and v', respectively. Indeed, if this were false, then either the ray $\vec{ov}$ would have to intersect b' in a segment disjoint from the segment ov, or the ray $\vec{ov'}$ would have to intersect b in a segment disjoint from ov'. However, either of these configurations would imply that γ crosses the lower boundary of either b or b', which, as we have just argued, is impossible. This establishes our claim. See Figure 9 for an illustration.

We choose some threshold parameter k, and consider the following two cases:

(a) $N_\sigma \le k$: Since the boundary of each of these N_σ sets intersects ∂a in at most $s/2$ vertices v of the type considered here, it follows that $\sigma \cap \partial a$ contains at most $sk/2$ such vertices. Summing this bound over all new slices σ with $N_\sigma \le k$, the overall number of vertices of U of this type in these slices is at most

$$O(k(\lambda_s(m_1) + m_2)\log\log n).$$

(b) $N_\sigma > k$: As argued above, the boundary of each of the N_σ sets of $\mathcal{C}_2$ that are incident to the vertices under consideration intersects at least k other such boundaries within the angular span of $W(\sigma)$. It follows that the arrangement $\mathcal{A}(\mathcal{C}_2)$ contains $\Omega(kN_\sigma)$ vertices at *level* at most k (i.e., vertices contained in at most k other sets of $\mathcal{C}_2$). On the other hand, arguing as in case (a), the number of vertices v of the above kind that are incident to $\sigma \cap \partial a$ is $\le sN_\sigma/2$. Hence, the number of these vertices is at most $O(N_\sigma/(kN_\sigma)) = O(1/k)$ times the number of vertices of $\mathcal{A}(\mathcal{C}_2)$ within $W(\sigma)$ at level at most k. Summing this inequality over all relevant slices σ, the overall number of such vertices is

$$O\left(\frac{1}{k}F_{\le k}(\mathcal{C}_2)\right),$$

where $F_{\le k}(\mathcal{C}_2)$ is the number of vertices of $\mathcal{A}(C_2)$ at level at most k. Using the Clarkson-Shor probabilistic analysis technique [3], we have $F_{\le k}(\mathcal{C}_2) = O(k^2 F(m_2/k))$, where $F(r)$ is the maximum number of vertices of the union of r compact convex α-fat sets, each pair of whose boundaries intersect at most s times (recall that α and s are assumed to be fixed parameters). By the induction hypothesis, we have $F(m_2/k) \le B(m_2/k)^{1+\varepsilon}$.

Hence, collecting all bounds, and taking into account regular vertices too (using Theorem 2.1), we obtain that the number of mixed vertices of U is at most

$$C\left(k(\lambda_s(m_1) + m_2)\log\log n + \frac{Bn^\varepsilon}{k^\varepsilon}m_2\right).$$

We now sum this bound over all nodes ξ and η of our segment trees and over the constant number of stabbing points used in each node, observing that the sums of the quantities m_1, m_2 are both $O(n\log^2 n)$. Hence the total number of vertices of the union is

$$C\left(k\lambda_s(n)\log^2 n\log\log n + \frac{\log^2 n}{k^\varepsilon}Bn^{1+\varepsilon}\right).$$

We now choose $k = n^{\varepsilon/2}$, and observe that the first term in this bound is at most $An^{1+\varepsilon}$, for a sufficiently large constant A, and the second term is at most

$$\frac{C\log^2 n}{n^{\varepsilon^2/2}} \cdot Bn^{1+\varepsilon}.$$

Hence this term is smaller than $\frac{1}{2}Bn^{1+\varepsilon}$, provided n is larger than some threshold $n_0(\varepsilon)$ that depends on ε, α and s. We now choose B so that (i) $B > 2A$ and (ii) $Bn^{1+\varepsilon}$ is an upper bound for the complexity of the union for each $n \le n_0(\varepsilon)$. The induction step is now complete.

This concludes the proof of Theorem 1.1. □

Acknowledgment

We wish to thank János Pach for helpful discussions concerning this problem. These discussions have led to the paper [10], on whose main result (Theorem 2.1) the current paper strongly depends.

References

[1] P. Agarwal, M. Katz and M. Sharir, Computing depth order and related problems, *Comput. Geom. Theory Appls.* 5 (1995), 187–206.

[2] B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, Algorithms for bichromatic line segment problems and polyhedral terrains, *Algorithmica* 11 (1994), 116–132.

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

[4] A. Efrat and M. Sharir, The complexity of the union of fat objects in the plane, *Proc. 13th ACM Symp. on Computational Geometry* (1997), 104–112.

[5] A. Efrat, M. Sharir and A. Ziv, Computing the smallest k-enclosing circle and related problems, *Comput. Geom. Theory Appls.* 4 (1994), 119–136.

[6] M. Katz, 3-D vertical ray shooting and 2-D point enclosure, range searching, and arc shooting amidst convex fat objects, to appear in *Comput. Geom. Theory Appls.* Also in a Tech. Report 2583, INRIA, Sophia-Antipolis, France, 1995.

[7] M. Katz, M. Overmars and M. Sharir, Efficient output sensitive hidden surface removal for objects with small union size, *Comput. Geom. Theory Appls.* 2 (1992), 223–234.

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

[9] J. Matoušek, J. Pach, M. Sharir, S. Sifrony and E. Welzl, Fat triangles determine linearly many holes, *SIAM J. Comput.* 23 (1994), 154–169.

[10] J. Pach and M. Sharir, On the boundary of the union of planar convex sets, to appear in *Discrete Comput. Geom.*

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

[12] F. van der Stappen, *Motion Planning amidst Fat Obstacles*, Ph.D. Dissertation, Utrecht University, 1994.

[13] M. van Kreveld, On fat partitioning, fat covering, and the union size of polygons, *Proc. 3rd Workshop Algorithms Data Struct.*, Lecture Notes in Computer Science, Vol. 709, Springer-Verlag, 1993, 452–463.