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

Alon Efrat[†] Micha Sharir[‡]

June 23, 2003

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 also been supported 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 $O(n)$ [7]. 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, an easy observation that has been noted in [9].

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 [3]), hidden surface removal in a collection of fat objects in 3-space [6], and point-enclosure queries in a collection of fat objects in the plane [5].

Remark: In an earlier version of this paper, the authors have proved the slightly better bound $O(\lambda_s(n))$ on the complexity of the union of $\mathcal{C}$, under the additional assumptions that all the regions in $\mathcal{C}$ are roughly of the same size and have bounded curvature. The proof for this special case is considerably simpler than the one given below. This result will be part of the final version of the paper, although it is not given in the present version.

2 Analysis

2.1 Touching and shattering vertices

In this subsection, we derive a general property of the union of planar sets, which we believe to be of independent interest.

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

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

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:

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

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

The *level* of a vertex of $\mathcal{A}(\mathcal{C})$ is the number of regions that contain it in their interior. Thus the vertices of U are exactly the vertices at level 0.

Let $T(\mathcal{C})$ (resp. $S(\mathcal{C})$) denote the number of touching (resp. shattering) vertices of U.

Theorem 2.1 *For any integer parameter $k < n/2$, we have*

$$T(\mathcal{C}) = O\left(n + k\mathbf{E}(T(\mathcal{R})) + k^2\mathbf{E}(S(\mathcal{R}))\right),$$

where $\mathcal{R}$ is a random sample of n/k regions of $\mathcal{C}$, and where $\mathbf{E}$ denotes expectation with respect to the choice of $\mathcal{R}$.

Proof: Fix a set $c \in \mathcal{C}$, and consider the circular sequence σ of vertices of $\mathcal{A}(\mathcal{C})$ in counterclockwise order along ∂c. Partition σ into contiguous subsequences $\sigma_1, \ldots, \sigma_m$, such that the arcs of ∂c between the first and last vertices of each σ_i (where we go in counterclockwise direction from the first to the last vertex) are precisely the connected components of $\partial c \cap int(U)$. See Figure 1 for an illustration. Let us denote by γ_i the arc corresponding to σ_i, for $i = 1, \ldots, m$, and let u_i, v_i denote the clockwise and counterclockwise endpoints of γ_i, respectively. Note that the overall number N of subsequences, over all $c \in \mathcal{C}$, is the number of vertices of U (each subsequence contributes two vertices and each vertex is counted twice).

Let k be the given 'threshold parameter'. We classify the subsequences σ_i (and the arcs γ_i) into the following three categories:

Short-and-touching sequences: Sequences σ_i with fewer than k vertices, all of which are touching.

Shattering sequences: Sequences σ_i that contain a shattering vertex among their first k elements.

Long sequences: Sequences σ_i with at least k elements, whose first k elements are all touching.

The arcs γ_i inherit the same classification from their corresponding sequences σ_i. We denote the overall number of short-and-touching (resp. shattering, long) subsequences, over all sets $c \in \mathcal{C}$, by N_{st} (resp. N_s, N_l). Clearly, $N = N_{st} + N_s + N_l$.

Lemma 2.2 $N_{st} = O(N_s + N_l + n)$.

Proof: We construct a (plane embedding of a) planar graph G as follows. The nodes of G are the counterclockwise endpoints v_i of shattering and of long arcs; in addition, if a region c has only short-and-touching arcs along its boundary, we add to G a node that lies somewhere on ∂c, but not on any of these arcs. (If this is impossible then $\partial c \subseteq int(U)$, and we can simply ignore c in what follows, since it does not contribute any vertex to U.) The number of nodes of G is thus $\leq N_s + N_l + n$.

Let w be a vertex of U incident to the boundaries of two sets $a, b \in \mathcal{C}$, such that w is an endpoint of a short-and-touching sequence on both ∂a and ∂b. Let v_a (resp. v_b) be the node of G nearest to w along ∂a (resp. ∂b) in clockwise direction. We then add to G an edge that connects v_a to v_b, and draw it by connecting v_a to w along ∂a in counterclockwise direction, and by connecting w to v_b along ∂b in clockwise direction. We refer to each of these two portions of the edge as a *half-edge*. We shift the resulting collection of edges slightly, to make sure that they do not overlap along the boundaries; the rule is that when several half-edges emerge from the same node v_a, their relative interiors are slightly shifted into the interior of the corresponding region a, so that the shorter the half-edge is, the closer it is to the boundary. See Figure 1 for an illustration.

This drawing of G may contain crossing pairs of edges, but we claim that any pair of edges cross an even number of times. To see this, consider the collection of half-edges of G, each connecting a node $v_a \in \partial a$, along the boundary of a, to a 'middle' vertex w, as defined above. We claim that a pair of half-edges is either disjoint or cross each other exactly twice. Indeed, by construction, half-edges along the boundary of the same a are drawn so that they do not cross at all. Let π and π' be two crossing half-edges, drawn along the boundaries of two respective distinct sets $c, c' \in \mathcal{C}$. By construction, all the sets that π crosses are such that their boundaries cross ∂c exactly twice, and π passes through (or, since it was perturbed, very near) those two intersection points, and similarly for π'. Hence ∂c and $\partial c'$ cross each other exactly twice, and the same holds for π, π'. (It is possible that one of these two points w is an endpoint of, say, π'. In this case π' must reach w from inside c, for otherwise it would not have crossed c at all, as is easily verified, contrary to our assumption. Since π has been perturbed into c, it follows that π and π' cross at w too.)

It is known [8, Corollary 3.1] that a graph that can be drawn in the plane so that every pair of edges cross an even number of times is planar. Hence G is planar. By construction, and by definition of touching vertices, any pair of nodes of G are connected by at most two edges, so Euler's formula is easily seen to imply that the number of edges of G is at most 6 times the number of its nodes. This completes the proof of the lemma. □

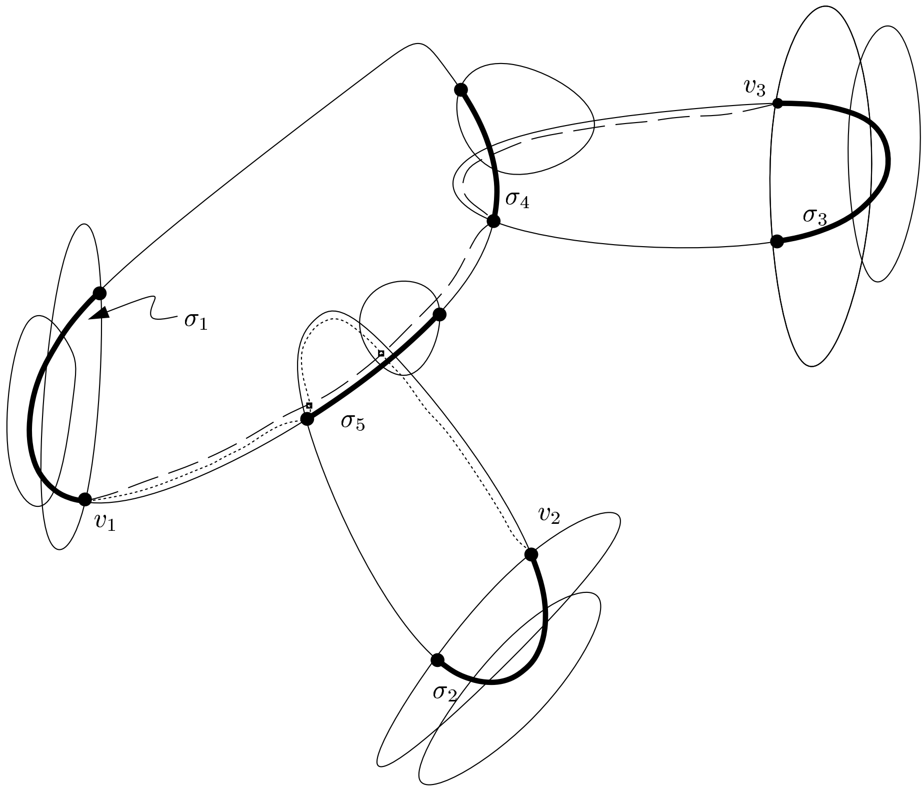

Figure 1: The proof of Theorem 2.1: σ_1, σ_2, and σ_3 are shattering sequences; σ_4 and σ_5 are short-and-touching sequences. Two edges of the graph G are drawn as dashed arcs.

Lemma 2.2 implies that $T(\mathcal{C}) = O(N_s + N_l + n)$. Indeed, the lemma bounds the number of touching vertices on ∂U that are endpoints of two short-and-touching arcs. Any other touching vertex of ∂U is an endpoint of a shattering or a long arc, so this upper bound caters to these vertices too. We charge each long sequence σ to the block of the first k of its vertices. Clearly, each such vertex can be charged at most twice (once along each boundary it is incident to), and they are all at level at most k. Hence, N_l is at most $2/k$ times the number of touching vertices at level at most k. Applying the probabilistic analysis technique of Clarkson and Shor [2] (see also [10]), we thus obtain

$$N_l = \frac{2}{k} \cdot O(k^2 \mathbf{E}(T(\mathcal{R}))) = O(k\mathbf{E}(T(\mathcal{R}))),$$

where $\mathcal{R}$ is a random sample of n/k regions from $\mathcal{C}$. Similarly, we charge each shattering sequence σ to the first shattering vertex that it contains. Again, such a vertex can be charged at most twice, and it lies at level at most k, so, arguing as above, we obtain $N_s = O(k^2\mathbf{E}(S(\mathcal{R})))$, where $\mathcal{R}$ is as above. This concludes the proof of Theorem 2.1. □

Remark: The proof technique of Theorem 2.1 can be applied to obtain an alternative simple proof of the result of [7], that if $\mathcal{A}(\mathcal{C})$ has no shattering vertices then the number of vertices of U is at most $6n - 12$ (for $n \geq 3$). For this, we define a graph G^* with n nodes, one node for each region in $\mathcal{C}$, where the node of region c is drawn as a point on ∂c that also lies on ∂U (if no such point exist, we can ignore c, as above). We draw an edge of G^* for each vertex w of U, incident to two boundaries ∂a, ∂b, by connecting, as above, the nodes representing a and b to w along (actually, slightly shifted away from) the respective boundaries. Arguing exactly as in the proof of Lemma 2.2, it is easily verified that every pair of edges of G^* cross an even number of times. Hence G^* is planar, and no pair of its nodes is connected by more than two edges, which implies, using Euler's formula as above, that the number of edges of G^* is at most $6n - 12$, as asserted.

2.2 Caps, inscribed fat polygons, and their properties

We now return to the case where $\mathcal{C}$ satisfies 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

$$\arcsin(\cos(\pi/t)/\alpha) > 2\pi/t,$$

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). P_c is defined to be the convex polygon whose vertices are $w_0, \ldots, w_{t-1}$. (Note that P_c may have fewer than t vertices if ∂c contains nonsmooth points whose tangent orientations span a sufficiently large interval.) The difference $c \setminus P_c$ is the union of at most t *caps* of

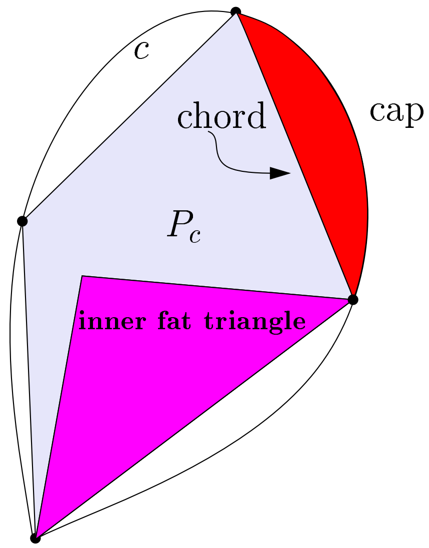

Figure 2: The inscribed polygon P_c and the corresponding caps; one inner fat triangle is also illustrated.

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 2.

Lemma 2.3 *The polygons P_c are α'-fat, for $\alpha' = \alpha / \cos(\pi/t)$.*

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 $r_2 \le \alpha r_1$. 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 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 $2\pi/t$. Let d denote the distance from the center O of D_1 to pq. It is easy to verify that $d \ge r_1 \cos(\pi/t)$. Indeed (see Figure 3), translate the tangents τ_p and τ_q so that they support $K \cap D_1$ at two respective points p', q'. The angle $p'Oq'$ is at most $2\pi/t$, so at least one of the angles between the perpendicular from O to $p'q'$ and $p'O$ or $q'O$ is at most π/t. Since both $|p'O|$ and $|q'O| = r_1$, the claim follows. This implies that the disk concentric with D_1 and having radius $r_1 \cos(\pi/t)$ 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 $r_2 \le \alpha' r_1$. 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 $\sin\beta = r_1/|Op| \ge r_1/r_2 \ge 1/\alpha'$. Similarly, the angle Oqp must also be at

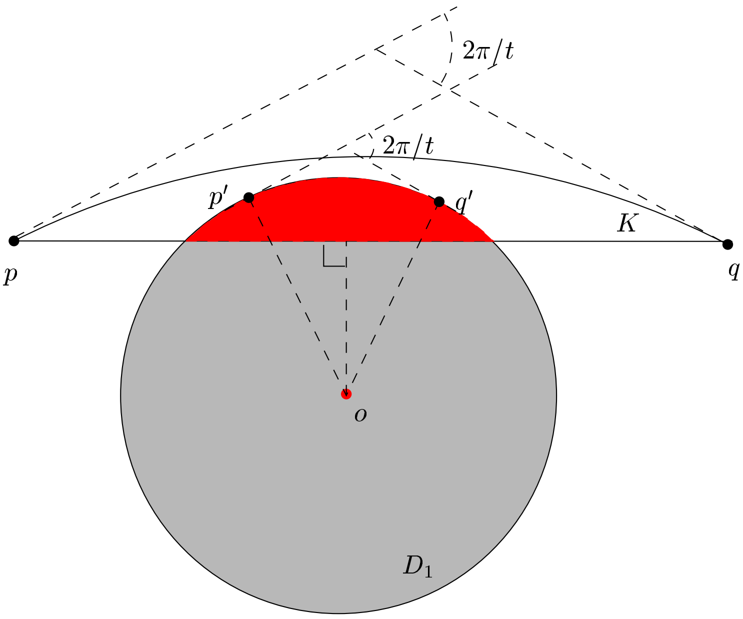

Figure 3: The proof of Lemma 2.3

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

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

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, $|\mathcal{T}| \le nt$. As an immediate consequence of [9], we have:

Lemma 2.4 *The union $U_{\mathcal{T}}$ of the triangles in $\mathcal{T}$ has $O(n \log\log n)$ vertices.*

Let v be a shattering 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 2.5 *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 4.

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 4.

Let θ_u^a, θ_w^a denote the orientations of the tangents to a at u and w, respectively, and let θ_u^b, θ_w^b 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 $(\theta_u^a, \theta_u^b, \theta_w^b, \theta_w^a)$, and they partition the circular range of orientations into four angular intervals that we denote by (θ_u^a, θ_u^b), (θ_u^b, θ_w^b), (θ_w^b, θ_w^a), and (θ_w^a, θ_u^a). 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 $\pi - 4\pi/t > 2\pi/3$.

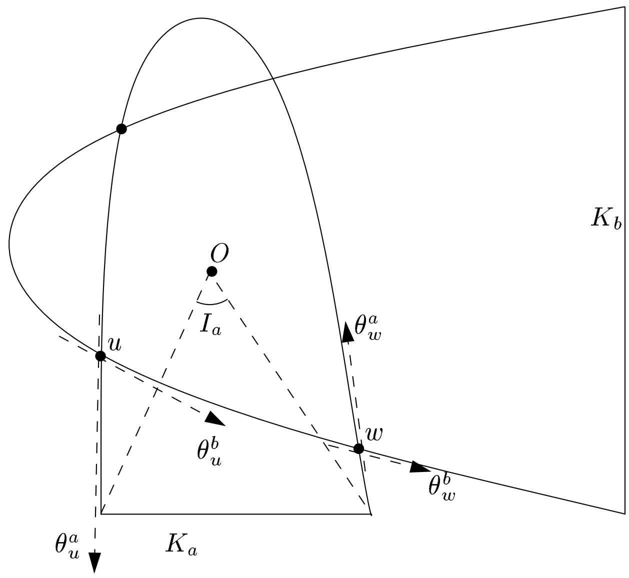

Figure 4: Two intersecting caps without a chordal intersection

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 2.6 *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 5(i)).

(ii) K_a contains a vertex of Δ_b (as in Figure 5(ii)).

(iii) Δ_b contains a vertex of K_a (as in Figure 5(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 5(iv.a)), or ∂a and ∂b intersect only twice (as in Figure 5(iv.b)).

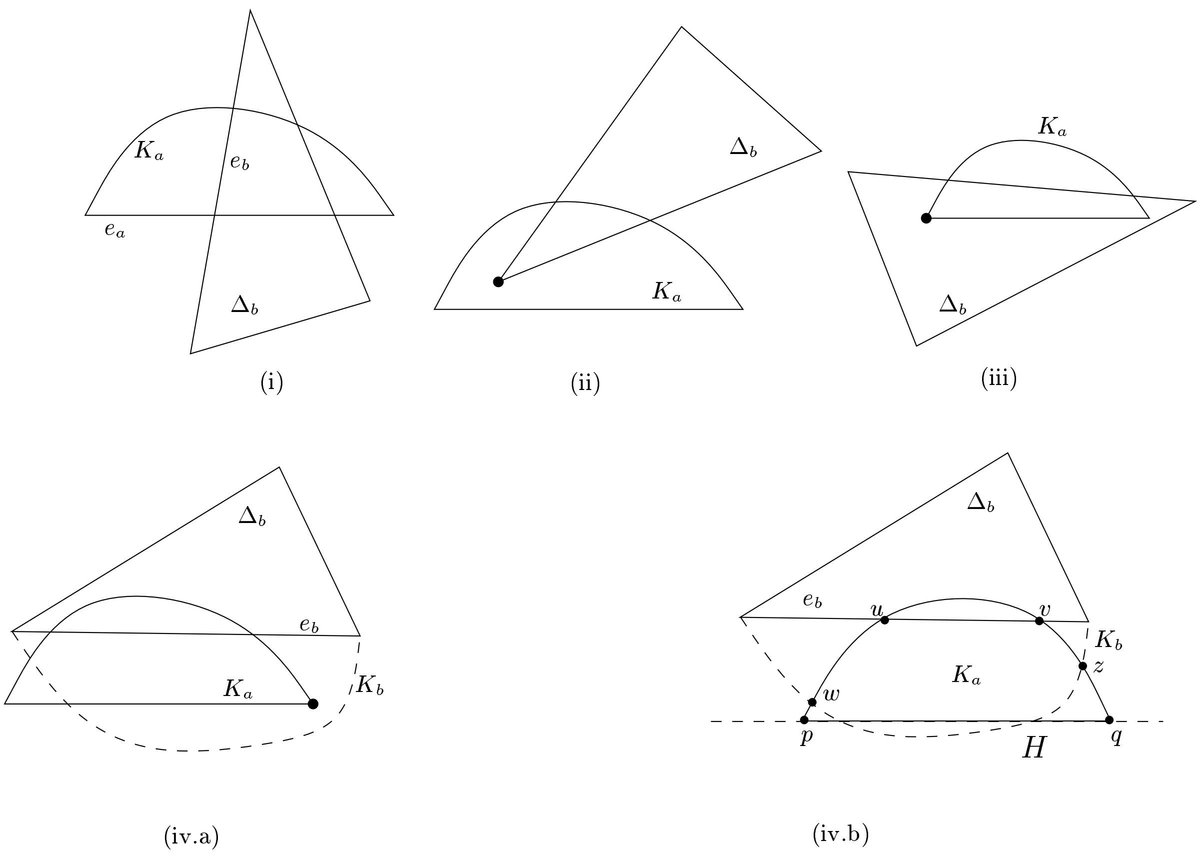

Figure 5: Illustrating the various cases in Lemma 2.6

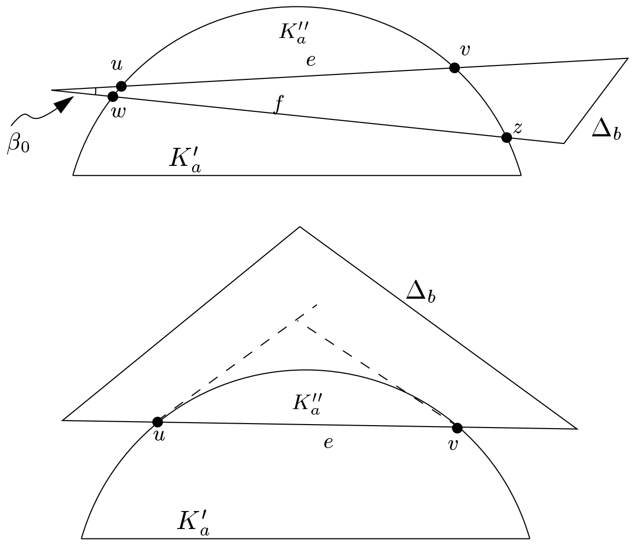

Figure 6: Two patterns of intersection of a cap K_a and an inner fat triangle Δ_b

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 $(\theta_0, \theta_0 + 2\pi/t)$. Clearly, the orientations of e_a and of e_b also lie in this range. Two cases can arise:

(1) Δ_b overlaps K'_a and is disjoint from K''_a (see Figure 6(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 $(\theta_0, \theta_0 + 2\pi/t)$, 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 6(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 $(\theta_0, \theta_0 + 2\pi/t)$, 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 5(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), ∂K_a and ∂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.) $\Box$

2.3 The proof of Theorem 1.1

The proof follows the technique used in the analysis of the complexity of lower envelopes of surfaces in higher dimensions and of related structures, as given in [4, 11].

Let $\mathcal{K} = \mathcal{K}(\mathcal{C})$ denote the collection of all caps of sets in $\mathcal{C}$, as defined above; recall that $|\mathcal{K}| \leq nt$. Let $U_{\mathcal{K}}$ denote the union of these caps. The vertices of U are also vertices of $U_{\mathcal{K}}$.

Let $K_a \subseteq a$ and $K_b \subseteq b$ be two caps of two (distinct) regions $a, b \in \mathcal{C}$, such that ∂K_a and ∂K_b intersect in at least one vertex that is shattering in $\mathcal{A}(\mathcal{C})$ (in other words, this is a vertex incident to ∂a and ∂b, and these boundaries cross at least four times; note also that a vertex can be shattering in $\mathcal{A}(\mathcal{C})$ and not in $\mathcal{A}(\mathcal{K})$ or vice versa). Put $R = K_a \cap K_b$. We call an arc of ∂R *marked* if it contains a vertex of $\mathcal{A}(\mathcal{K})$ that lies on the chord of some cap, such that condition (iv) of Lemma 2.6 does not hold for that vertex (or, rather, for the cap and the triangle on whose boundaries the vertex lies), and *unmarked* otherwise; see Figure 7. We will refer to vertices of $\mathcal{A}(\mathcal{K})$ that lie on some chord as *chordal* vertices. Chordal vertices that satisfy condition (iv) of Lemma 2.6 will be called *special chordal vertices*, and all the other chordal vertices will be called *standard chordal vertices.*

Lemma 2.5 and Lemma 2.6 imply that at least one arc of ∂R is marked. Actually, they imply that one of the vertices of R is a standard chordal vertex. Indeed, Lemma 2.5 implies that R has at least one chordal vertex. If all the chordal vertices

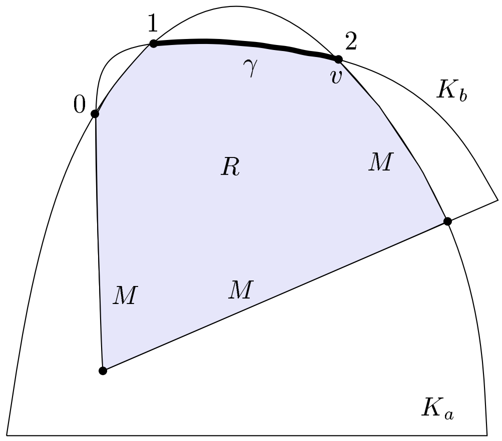

Figure 7: The region $R = K_a \cap K_b$; the marked arcs are labeled by M; ∂R has three shattering vertices, with indices 0,1,2, as shown

of R are special then each of the chords of K_a, K_b contains zero or two such vertices, and, arguing as in the proof of Lemma 2.6, one can show that ∂a and ∂b intersect in exactly two vertices, contradicting the assumption that ∂K_a and ∂K_b intersect in at least one shattering vertex in $\mathcal{A}(\mathcal{C})$. Note also that if ∂R has any special chordal vertex, then it must also have an endpoint of a chord as a vertex (see Lemma 2.6).

We define the *index* of a vertex w of $a \cap b$ that lies on ∂R to be the number of unmarked arcs of ∂R that we encounter before hitting a marked arc, as we traverse this boundary from w in counterclockwise direction. However, if ∂R has a special chordal vertex, it can have at most one vertex of $\partial a \cap \partial b$, as follows easily from the analysis of the proof of Lemma 2.6, and we define the index of that vertex to be 0. The index of a vertex is an integer between 0 and $s-1$.

We also define the *level* of a vertex v of the arrangement $\mathcal{A}(\mathcal{K})$ to be the number of caps of $\mathcal{K}$ containing v in their interior. Clearly, vertices at level 0 are exactly the vertices of $\partial U_{\mathcal{K}}$.

We define the following quantities:

- $T(\mathcal{C})$ is the number of touching vertices in $\mathcal{A}(\mathcal{C})$ that lie on ∂U.
- $S(\mathcal{C})$ is the number of shattering vertices in $\mathcal{A}(\mathcal{C})$ that lie on ∂U.
- $S^{(j)}(\mathcal{C})$, for $j = 0, \ldots, s-1$, is the number of these shattering vertices whose index is at most j (so $S(\mathcal{C}) = S^{(s-1)}(\mathcal{C})$).
- $Q^*(\mathcal{C})$ is the number of special chordal vertices that lie at level 0 in the arrangement $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$ of the caps and inner fat triangles of the regions in $\mathcal{C}$.[2]

[2]One might be tempted to think that these vertices lie on the boundary of the union of these

- $Q_0(\mathcal{C})$ is the number of standard chordal vertices that lie at level 0 in $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$.
- $T(n)$ is the maximum of $T(\mathcal{C})$, over all collections $\mathcal{C}$ of n sets satisfying (i) and (ii) (with fixed α and s).
- $S(n)$ is the maximum of $S(\mathcal{C})$, over all collections $\mathcal{C}$ as above.
- $S^{(j)}(n)$ is the maximum of $S^{(j)}(\mathcal{C})$, for $j = 0, \ldots, s-1$, over all collections $\mathcal{C}$ as above.
- $Q^*(n)$ is the maximum of $Q^*(\mathcal{C})$, over all collections $\mathcal{C}$ as above.
- $Q_0(n)$ is the maximum of $Q_0(\mathcal{C})$, over all collections $\mathcal{C}$ as above.

We will derive a (somewhat complex) system of recurrence relationships for the above quantities. Each of these recurrences involves a 'threshold parameter' $k < n/2$, which is arbitrary, and we will choose a different value of k for each recurrence, in a manner detailed below.

First, using Theorem 2.1, we have:

$$T(n) \le c\left(n + kT(n/k) + k^2 S(n/k)\right), \tag{1}$$

for some constant c (for simplicity, we will use the same constant in all the recurrences).

We next estimate $S^{(j)}(n)$, for $j = 0, \ldots, s-1$. Let v be a shattering vertex of $\mathcal{A}(\mathcal{C})$ that lies on ∂U, incident to two boundaries ∂a, ∂b, and contained in two respective caps $K_a \subseteq a$ and $K_b \subseteq b$, whose index is at most j. Let $R = K_a \cap K_b$, and let γ denote the arc of ∂R incident to v and lying counterclockwise to it; see Figure 7. We fix a threshold parameter k, trace γ from v, and examine the sequence σ of vertices of $\mathcal{A}(\mathcal{K})$ that we encounter. Several cases can arise:

(a) σ contains at least k vertices, and none of the first k vertices of σ is a standard chordal vertex. (Note that σ may contain a standard chordal vertex only if the index of v is zero.) In this case we charge v to the block of the first k vertices of σ. Several important observations need to be made:

 (i) Each of these charged vertices lies at level $\le k$ (in $\mathcal{A}(\mathcal{K})$, or in $\mathcal{A}(\mathcal{C})$ if it is a vertex of this latter arrangement).

 (ii) Each charged vertex is charged at most twice, once along each of the boundaries containing it.

caps and triangles, but this is not the case, since each chordal vertex z lies in the interior of the union of the cap and triangle that share the chord that contains z. This is why we use the more careful notion of level 0.

(iii) The portion of γ that contains σ may lie inside some inner fat triangles, but it does not cross any non-chordal edge of any such triangle (this follows from condition (iv) of Lemma 2.6).

(iv) Each of the charged special chordal vertices lies at level $\leq k$ in $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$ (this is a consequence of observation (iii)).

It follows that the overall number of shattering vertices of $\mathcal{A}(\mathcal{C})$ that lie on ∂U and fall into this category is at most $2/k$ times the number of touching and shattering vertices of $\mathcal{A}(\mathcal{C})$ that lie at level at most k (in either arrangement) plus $2/k$ times the number of special chordal vertices that lie at level at most k in $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$ (this is a consequence of observation (iii)). Using again the probabilistic analysis technique of Clarkson and Shor [2], this upper bound is at most $(2/k) \cdot O(k^2)$ times the expected number of touching and shattering vertices of the union of a random sample of n/k regions of $\mathcal{C}$, plus $(2/k) \cdot O(k^2)$ times the expected number of special chordal vertices at level 0 in the arrangement of the caps and inner fat triangles of a random sample of n/k regions of $\mathcal{C}$. In other words, the number of vertices v of the present type is $O(kT(n/k) + kS(n/k) + kQ^*(n/k))$.

(b) The first k (or all, if σ is shorter) vertices in σ include at least one standard chordal vertex w (in this case the index of v must be zero). We charge v to the first such w. It is easily verified that w can be charged in this manner only once. Since w is at level at most k in $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$ (see observation (iv)), another application of the Clarkson-Shor technique implies that the number of vertices v of this kind is at most $O(k^2 Q_0(n/k))$.

(c) σ contains fewer than k vertices of $\mathcal{A}(\mathcal{K})$, none of which is standard chordal. In this case we consider the other endpoint w of γ. If w is also a shattering vertex in $\mathcal{A}(\mathcal{C})$, then its index is at most $j-1$, and it lies at level at most k in $\mathcal{A}(\mathcal{C})$. In this case we charge v to w. Otherwise, w is a special chordal vertex, in which case, as observed above, ∂R must contain an endpoint z of the chord of K_a or of K_b. Moreover, suppose, without loss of generality, that $\gamma \subseteq \partial a$. Then the proof of Lemma 2.6 is easily seen to imply that z is an endpoint of the chord of K_a, and that the entire counterclockwise portion of ∂a from v to z is lies in the interior of $U_{\mathcal{K}}$ (see Figure 5(iv.b)). In this case we charge v to z. The argument just given implies that z can be charged at most once in this manner, and the number of such points z is at most $nt = O(n)$. Applying the Clarkson-Shor technique again to the former type of charging, we conclude that the number of vertices v of this kind is $O(n + k^2 S^{(j-1)}(n/k))$.

Thus, summing up these bounds, we obtain the following system of recurrences:

$$S^{(j)}(n) \leq c \left(n + kT(n/k) + kS(n/k) + kQ^*(n/k) + k^2 S^{(j-1)}(n/k) + k^2 Q_0(n/k)\right) , \quad (2)$$

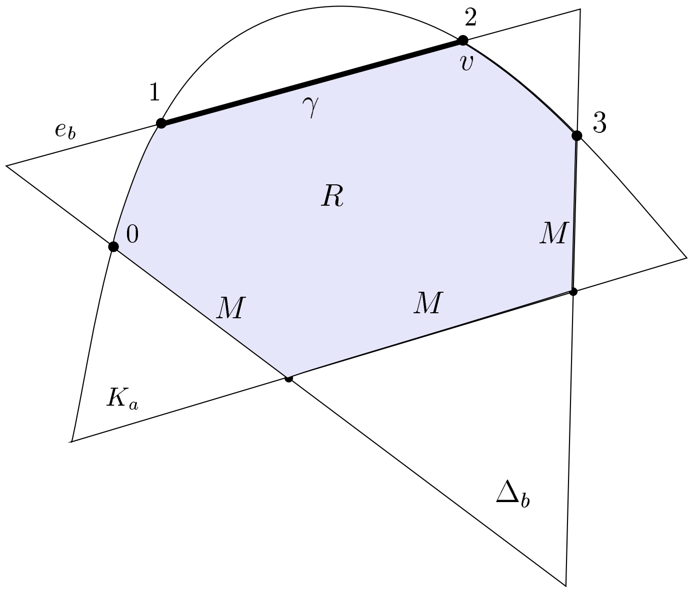

Figure 8: The region $R = K_a \cap \Delta_b$; the marked arcs are labeled by M; ∂R has four semi-sharp vertices, with indices 0,1,2,3, as shown

for $j = 0, \ldots, s-1$, and for some constant c that depends on α (for $j = 0$, we put $S^{(-1)} = 0$ in the right-hand side).

Next we estimate $Q_0(\mathcal{C})$, using an analysis similar to the one just presented. Recall that we are counting standard chordal vertices at level 0 in $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$. (Formally, we can also take each pair of a cap and a triangle with a common chord, and shift them slightly away from each other, so that each of the standard chordal vertices that lie on this chord and are counted in $Q(\mathcal{C})$ is split into two vertices, both of which lie on the boundary of the union of $\mathcal{K} \cup \mathcal{T}$.)

We analyze the number of these chordal vertices by considering them as vertices in $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$, each lying on the boundary of the intersection of an inner fat triangle and a cap. So let v be a standard chordal vertex, lying on the boundary of a cap $K_a \subseteq a$ and on the chord of an inner fat triangle $\Delta_b \subseteq b$, for two (distinct) sets $a, b \in \mathcal{C}$. By definition, K_a and Δ_b satisfy one of the conditions (i)–(iii) of Lemma 2.6, which means that the boundary of the intersection $R = K_a \cap \Delta_b$ has at least one vertex that is a vertex of $\mathcal{A}(\mathcal{T})$, and none of its vertices is a special chordal vertex (see Figure 8). Observe that ∂R consists of at most six arcs, since each edge of Δ_b can intersect ∂K_a at most twice. We call an arc of ∂R *marked* if it contains a vertex of $\mathcal{A}(\mathcal{T})$, and *unmarked* otherwise. We will refer to these vertices as *sharp* vertices. As just argued, at least one arc of ∂R is marked. We define the index of any non-sharp vertex w of R (including non-chordal vertices as well) to be the number of unmarked arcs of ∂R that we encounter before hitting a marked arc, as we traverse this boundary from w in counterclockwise direction. The index of a vertex is an integer between 0 and 5. As just mentioned, not all the vertices that we encounter during this traversal need be chordal; each non-chordal and non-sharp vertex lies on ∂a and on some edge (other than the chord) of Δ_b. We refer to all such vertices as *semi-sharp*; clearly, chordal vertices are also semi-sharp.

We introduce more quantities that we want to bound:

- $Q(\mathcal{C})$ is the number of semi-sharp vertices on $\partial U_{\mathcal{K}\cup\mathcal{T}}$ (perturbed as above), excluding special chordal vertices. Clearly, $Q_0(\mathcal{C}) \le Q(\mathcal{C})$.
- $Q^{(j)}(\mathcal{C})$, for $j = 0, \ldots, 5$, is the number of semi-sharp vertices on $\partial U_{\mathcal{K}\cup\mathcal{T}}$ (excluding special chordal vertices), whose index is at most j (so $Q(\mathcal{C}) = Q^{(5)}(\mathcal{C})$).
- $Q(n)$ the maximum of $Q(\mathcal{C})$, over all collections $\mathcal{C}$ of n regions satisfying (i) and (ii) (with fixed α and s).
- $Q^{(j)}(n)$ is the maximum of $Q^{(j)}(\mathcal{C})$, over all such collections $\mathcal{C}$, for $j = 0, \ldots, 5$.

Let v be a semi-sharp vertex of $\partial U_{\mathcal{K}\cup\mathcal{T}}$, incident to the boundary of a cap $K_a \subseteq a$ (and to ∂a itself) and to the boundary of an inner fat triangle $\Delta_b \subseteq b$, for two distinct sets $a, b \in \mathcal{C}$, whose index is at most j. Let $R = K_a \cap \Delta_b$, and let γ denote the arc of ∂R incident to v and lying counterclockwise to it. We fix some threshold parameter k, trace γ from v, and examine the sequence σ of vertices of $\mathcal{A}(\mathcal{K}\cup\mathcal{T})$ that we encounter. Several cases can arise:

(a) σ contains at least k vertices, and none of the first k vertices of σ is sharp. (Note that σ may contain a sharp vertex only if the index of v is zero.) In this case we charge v to the block of the first k vertices of σ. As above, we have the important observations that (i) each of these vertices lies at level $\le k$ in $\mathcal{A}(\mathcal{K}\cup\mathcal{T})$, and (ii) each such vertex is charged at most twice, once along each of the boundaries containing it. It follows that the overall number of semi-sharp vertices of $\partial U_{\mathcal{K}\cup\mathcal{T}}$ that fall into this category is at most $(2/k)$ times the number of touching and shattering vertices of $\mathcal{A}(C)$ at level at most k, and of semi-sharp and special chordal vertices of $\mathcal{A}(\mathcal{K}\cup\mathcal{T})$ at level at most k. As above, the probabilistic analysis technique of Clarkson and Shor implies that this upper bound is at most $(2/k)\cdot O(k^2)$ times the expected number of touching and shattering vertices of the union of a random sample of n/k regions of $\mathcal{C}$, plus $(2/k)\cdot O(k^2)$ times the expected number of semi-sharp and special chordal vertices of the union of the caps and triangles of a random sample of n/k regions of $\mathcal{C}$. Hence the number of semi-sharp vertices of $\partial U_{\mathcal{K}\cup\mathcal{T}}$ of this kind is $O(kT(n/k) + kS(n/k) + kQ(n/k) + kQ^*(n/k))$.

(b) The first k (or all, if σ is shorter) vertices in σ include at least one sharp vertex w (in this case the index of v must be zero). We charge v to the first such w. As above, it is easily seen that w can be charged in this manner at most twice, and that it lies at level at most k in $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$. Hence, using Lemma 2.4 and the Clarkson-Shor analysis technique, the number of vertices v of this kind is easily seen to be $O(k^2 \cdot \frac{n}{k} \log\log \frac{n}{k}) = O(nk \log\log n)$.

(c) σ contains fewer than k vertices of $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$, none of which is sharp. In this case we charge v to the other endpoint w of γ. Clearly, w is also a semi-sharp vertex (as noted above, no vertex of R is special chordal), whose index is at most $j-1$, and it lies at level at most k in $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$. Applying the Clarkson-Shor technique again, we conclude that the number of vertices v of this kind is $O(k^2 Q^{(j-1)}(n/k))$.

Thus, summing up these bounds, we obtain:

$$Q^{(j)}(n) \le c\left(nk \log\log n + kT(n/k) + kS(n/k) + kQ(n/k) + kQ^*(n/k) + k^2 Q^{(j-1)}(n/k)\right), \tag{3}$$

for $j = 0, \ldots, 5$, and for some constant $c > 0$ that depends on α (for $j = 0$, we put $Q^{(-1)} = 0$ in the right-hand side).

Finally, we estimate $Q^*(\mathcal{C})$, that is, the number of special chordal vertices on $\partial U_{\mathcal{K} \cup \mathcal{T}}$ (perturbed as above). We do this by applying a variant of the analysis in the proof of Theorem 2.1. Note that the special chordal vertices are touching vertices in $\mathcal{A}(\mathcal{K} \cup \mathcal{T})$, but not all such touching vertices are necessarily special chordal. We apply the same graph construction as in the proof of Theorem 2.1, except that instead of touching vertices we consider only the subset of special chordal vertices. Thus, for example, a subsequence σ_i (as in the proof of Theorem 2.1) that contains (among its first k members) a touching vertex that is not special chordal will be treated as a 'shattering' subsequence, and will correspond to a node of the graph (rather than to an edge of it), and edges are induced only by short sequences that consist exclusively of special chordal vertices. We leave it to the reader to verify that the proof remains valid under this modification, and that the conclusion now is that

$$Q^*(\mathcal{C}) = O\left(n + k\mathbf{E}(Q^*(\mathcal{R})) + k^2(\mathbf{E}(T(\mathcal{R})) + \mathbf{E}(S(\mathcal{R})) + \mathbf{E}(Q(\mathcal{R})))\right),$$

where $\mathcal{R}$ is a random sample of n/k regions of $\mathcal{C}$, as above. In other words, we have:

$$Q^*(n) \le c\left(n + kQ^*(n/k) + k^2 T(n/k) + k^2 S(n/k) + k^2 Q(n/k)\right), \tag{4}$$

for some constant c, as above.

Following the analysis in [4, 11], the solution of the combined recurrences (1), (2), (3), and (4), with appropriate choice of the threshold parameters k, can be shown to be

$$\begin{aligned} T(n) &= O(n^{1+\varepsilon}) \\ S(n) &= O(n^{1+\varepsilon}) \\ Q(n) &= O(n^{1+\varepsilon}) \\ Q^*(n) &= O(n^{1+\varepsilon}), \end{aligned}$$

for any $\varepsilon > 0$, where the constants of proportionality depend on ε, α, and s. Since these recurrences are somewhat more involved than those in [4, 11], we include, in the Appendix below, a proof of these bounds, for the sake of completeness. This concludes the proof of Theorem 1.1. □

Acknowledgment

We wish to thank János Pach for helpful discussions concerning this problem. His comments were especially helpful for the analysis given in Subsection 2.1.

References

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

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

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

[4] D. Halperin and M. Sharir, Almost tight upper bounds for the single cell and zone problems in three dimensions, *Discrete Comput. Geom.* 14 (1995), 385–410.

[5] 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.

[6] 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.

[7] 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.

[8] L. Lovász, J. Pach and M. Szegedy, On Conway's thrackle conjecture, *Proc. 11th ACM Symp. on Computational Geometry*, 1995, 147–151.

[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] M. Sharir, On k-sets in arrangements of curves and surfaces, *Discrete Comput. Geom.* 6 (1991), 593–613.

[11] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* 12 (1994), 327–345.

[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.

Appendix: Solving the Recurrences

In this appendix we prove that the solution of the above system of recurrences is near-linear. Recall that the recurrences are:

$$\begin{aligned}
T(n) &\leq c\left(n + kT(n/k) + k^2 S(n/k)\right) \\
S^{(j)}(n) &\leq c\left(kT(n/k) + kS(n/k) + kQ^*(n/k) + k^2 S^{(j-1)}(n/k) + k^2 Q(n/k)\right) \\
Q^{(j)}(n) &\leq c\,(nk \log\log n + kT(n/k) + kS(n/k) + kQ(n/k) + kQ^*(n/k) \\
&\qquad + k^2 Q^{(j-1)}(n/k)\big) \\
Q^*(n) &\leq c\left(n + kQ^*(n/k) + k^2 T(n/k) + k^2 S(n/k) + k^2 Q(n/k)\right).
\end{aligned}$$

Before deriving the formal solution, here is an intuitive explanation of the analysis. The right-hand sides of these recurrences include three kinds of terms:

(i) 'overhead', non-recursive terms that are linear or near-linear in n,

(ii) recursive terms with coefficients of the form $O(k)$, and

(iii) recursive terms with coefficients of the form $O(k^2)$.

If it were not for the terms of the third kind, the recurrences would trivially solve to $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$. Terms of the third kind are 'dangerous', because they 'suggest' a quadratic solution. Fortunately, though, there is a strict hierarchy between the various functions appearing in the recurrences, such that any term with a coefficient $O(k^2)$ involves a function that is lower in the hierarchy than the function appearing in the left-hand side. This hierarchy is

$$Q^* \to T \to S = S^{(s-1)} \to S^{(s-2)} \to \cdots \to S^{(0)} \to Q = Q^{(5)} \to Q^{(4)} \to \cdots \to Q^{(0)}.$$

There are $m = s + 8$ functions in this hierarchy, and we assign to each function F a *serial number* $i(F)$ in the hierarchy, so that the serial number of Q^* is $m - 1$ and that of $Q^{(0)}$ is 0. We exploit this hierarchy by choosing different k's in different recurrences, so that the k's chosen for recurrences whose left-hand-side functions are lower in the hierarchy are much larger than those chosen for functions higher in the hierarchy. In this way, the effect of the coefficients $O(k^2)$ can be made negligible, making the overall solution near-linear.

In more detail, we fix $\varepsilon > 0$. In the recurrence for the function F whose serial number is $i = i(F)$, for $i = 0, 1, \ldots, m-1$, we choose $k = k_F = k_0^{\varepsilon^i}$, for some sufficiently large k_0 that we will choose later. We claim that the solution of these recurrences is $F(n) \le A_F n^{1+\varepsilon}$, where F is any of the m functions Q^*, T, $S^{(j)}$, and $Q^{(j)}$, and A_F is the constant

$$A_F = A(6c)^{i(F)} k_0^{\varepsilon - \varepsilon^{i(F)+1}},$$

for some sufficiently large constant A that will be determined later.

We prove these upper bounds by induction on n. First, by choosing A to be sufficiently large, we may assume that these bounds hold for all the functions and for any $n \le n_0$, for some sufficiently large n_0 (that will be fixed below).

We put $A_{\max} = \max A_F$, over all the functions F in the recurrences. Note that A_F is a monotone increasing function of the serial number of F, so $A_{\max} = A_{Q^*} = A(6c)^{m-1} k_0^{\varepsilon - \varepsilon^m}$.

Each of the recurrences has the form

$$F(n) \le c\left(N(n) + k(G_{j_1}(n/k) + \cdots + G_{j_q}(n/k)) + k^2(F_{i_1}(n/k) + \cdots + F_{i_p}(n/k))\right),$$

for $k = k_F$, as defined above, where $N(n)$ is a near-linear overhead term (which may also depend on k), and the serial numbers of the functions $F_{i_1}, \ldots, F_{i_p}$ are smaller than that of F. Moreover, in all the recurrences we have $1 + q + p \le 6$. Using the induction hypothesis, we need to show that, for $n > n_0$,

$$\frac{cN(n)}{n^{1+\varepsilon}} + \frac{c(A_{G_{j_1}} + \cdots + A_{G_{j_q}})}{k_F^\varepsilon} + \frac{ck_F(A_{F_{i_1}} + \cdots + A_{F_{i_p}})}{k_F^\varepsilon} \le A_F.$$

Let G be the function immediately following F in the hierarchy (so $i(G) = i(F) - 1$). The monotonicity of the coefficients A_F implies that it suffices to show that

$$\frac{cN(n)}{n^{1+\varepsilon}} + \frac{cqA_{\max}}{k_F^\varepsilon} + \frac{cpk_F A_G}{k_F^\varepsilon} \le A_F.$$

Let $i = i(F)$. Then $k_F = k_0^{\varepsilon^i}$, $A_F = A(6c)^i k_0^{\varepsilon - \varepsilon^{i+1}}$, and $A_{\max} = A(6c)^{m-1} k_0^{\varepsilon - \varepsilon^m}$. We thus need to show that

$$\frac{cN(n)}{n^{1+\varepsilon}} + \frac{cqA(6c)^{m-1} k_0^{\varepsilon - \varepsilon^m}}{k_0^{\varepsilon^{i+1}}} + \frac{cpk_0^{\varepsilon^i} A(6c)^{i-1} k_0^{\varepsilon - \varepsilon^i}}{k_0^{\varepsilon^{i+1}}} \le A(6c)^i k_0^{\varepsilon - \varepsilon^{i+1}}.$$

Write the left-hand side as $L_1 + L_2 + L_3$, and the right-hand side as R. Note that R is minimized when $i = 0$, in which case it is equal to A. The term L_1 is a decreasing function of n, for $n > n_0$, assuming that n_0 is sufficiently large, and we choose A so that it satisfies

$$6\frac{cN(n_0)}{n_0^{1+\varepsilon}} \le A,$$

implying that $L_1 \leq R/6$ (for $n > n_0$).

Concerning L_2, we have

$$\frac{L_2}{R} = \frac{cq(6c)^{m-1-i}}{k_0^{\varepsilon^m}} \leq \frac{cq(6c)^{m-1}}{k_0^{\varepsilon^m}}.$$

We choose k_0 so that

$$(6c)^m \leq k_0^{\varepsilon^m},$$

which implies that $L_2 \leq qR/6$.

Finally, it is straightforward to check that $L_3/R = p/6$, or $L_3 \leq pR/6$.

Adding up these three inequalities, we have $L_1 + L_2 + L_3 \leq (1 + q + p)R/6 \leq R$, which completes the inductive proof. (To complete the argument, we need to fix n_0, we choose it large enough so that it is greater than $2k_0$ and so that all the overhead terms $N(n)$ are such that $N(n)/n^{1+\varepsilon}$ are decreasing functions of n.)