The Complexity of the Union of (α, β)-Covered Objects*

Alon Efrat†

January 28, 2000

Abstract

An (α, β)-covered object is a simply connected planar region c with the property that for each point $p \in \partial c$ there exists a triangle contained in c and having p as a vertex, such that all its angles are at least α and all its edges are at least $\beta \cdot diam(c)$-long. This notion extends that of fat convex objects. We show that the combinatorial complexity of the union of n (α, β)-covered objects of 'constant description complexity' is $O(\lambda_{s+2}(n) \log^2 n \log\log n)$, where s is the maximum number of intersections between the boundaries of any pair of the given objects.

1 Introduction

A planar object c is (α, β)-*covered* if the following conditions are satisfied.

1. c is simply-connected;

2. For each point $p \in \partial c$ we can place a triangle Δ fully inside c, such that p is a vertex of Δ, each angle of Δ is at least α, and the length of each edge of Δ is at least $\beta \cdot diam(c)$. We call such a triangle Δ a *good triangle* for c.

The notion of (α, β)-covered objects generalizes the notion of convex fat objects. A planar convex object c is α-*fat* if the ratio between the radii of the balls s^+ and s^- is at most α, where s^+ is the smallest ball containing c and s^- is a largest ball that is contained in c. It is easy to show that an α-fat convex object is an (α', β')-covered object, for appropriate constants α', β' that depend on α.

In this paper we will also make the additional assumption that all the objects under consideration have *constant description complexity*, meaning that each object is a semialgebraic set defined by a constant number of polynomial equalities and inequalities of constant maximum degree.

The goal of this paper is to obtain sharp bounds for the combinatorial complexity of the union of a collection $\mathcal{C}$ of n (α, β)-covered objects of constant description complexity, for constant parameters $\alpha, \beta > 0$.

*Work by Alon Efrat has been supported by Rothschild Fellowship and by NSF grant CCR-9623851.
†Computer Science Department, Stanford University, *alon@graphics.stanford.edu*

There are not too many results of this kind. If $\mathcal{C}$ is a collection of α-fat triangles[1], then the complexity of $\bigcup\mathcal{C}$ is $O(n \log\log n)$ (with the constant of proportionality depending on α) [12], and this bound improves to $O(n)$ if the triangles are nearly of the same size [2] or are infinite wedges. See also [11] 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 $\bigcup\mathcal{C}$ is $O(n)$ [10]. Of course, if we drop the fatness condition, the complexity of $\bigcup\mathcal{C}$ can be $\Omega(n^2)$, even for the case of (non-fat) triangles. Even for fat convex objects, some bound on the description complexity of each object must be assumed, or else the complexity of the union might be arbitrarily large.

In [4] it was shown that if $\mathcal{C}$ is a collection of n convex α-fat objects of constant description complexity, then the complexity of $\bigcup\mathcal{C}$ is $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε, on the fatness parameter and on the maximum description complexity of the given objects. In an attempt to remove the convexity restriction, it was shown in [5] that if $\mathcal{C}$ is a collection of n κ*-curved* objects of constant description complexity, then the complexity of $\bigcup\mathcal{C}$ is $O(\lambda_{s'}(n)\log^2 n)$, where s' is a constant that depends on κ and on the description complexity of the objects. A planar object c is κ-curved (for a parameter κ) if each point p on its boundary is contained in some disk $B \subseteq c$ whose radius is at least $\kappa \cdot diam(c)$. However, the class of κ-curved objects is rather restricted. For example, an α-fat triangle is not a κ-curved object for any κ. However, the notion of (α, β)-covered objects clearly generalizes the notion of κ-curved objects (as well as that of fat convex objects).

Let $\mathcal{C}$ be a collection of n (α, β)-covered objects of constant description complexity in general position. This implies that the boundaries of each pair of objects of $\mathcal{C}$ have at most some constant number, s, of intersection points, and we may assume that s also bounds the number of points at which the boundary of an object in $\mathcal{C}$ is not C^1 or C^2, as well as the number of inflection points and locally x- and y-extremal points of any such boundary.

The main result of this paper is:

Theorem 1.1 *Under the above assumptions, the combinatorial complexity of the union of* $\mathcal{C}$ *is* $O(\lambda_{s+2}(n)\log^2 n \log\log n)$.

The proof of Theorem 1.1 is given in the following sections. In is worth mentioning that if all objects of $\mathcal{C}$ are roughly of the same size, then the bound of Theorem 1.1 can be improved to $O(\lambda_{s+2}(n))$, see Remark 3.4 for further discussion.

Theorem 1.1, as well as the previous works cited above, contribute to 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 [7]), hidden surface removal in a collection of fat objects in 3-space [9], point-enclosure queries in a collection of fat objects in the plane [8], and more; See [16] for more applications, and other definitions of fat non-convex objects. Theorem 1.1 both extends these results to the more general class of (α, β)-covered objects, and slightly improves the corresponding complexity bounds.

The contributions of this paper are thus (a) the introduction of the new class of 'fat' non-convex objects (namely (α, β)-covered objects), which, as we believe, captures the input

[1]For triangles, there is an equivalent definition of fatness that requires all angles to be at least some fixed constant α_0; in [12], this is called α_0-fatness.

data in most realistic scenes; (b) presenting a sharper bound on the union complexity than the bounds obtained in [4] (bringing them to within a polylogarithmic factor off the actual complexity); and (c) the proof technique, which is much simpler than the analysis given in [4].

2 Preliminaries

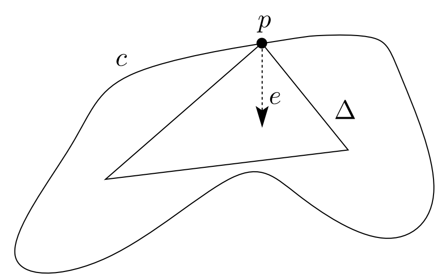

Figure 1: The point p is $3\pi/2$-oriented.

Let $\mathcal{C}$ be a collection of n (α,β)-covered objects, as in the introduction. Let $c \in \mathcal{C}$, and let p be a point on ∂c. We say that p is *θ-oriented* if there is a good triangle Δ for c with p as a vertex, such that the ray e emerging from p at orientation θ intersects the interior of Δ. In this case we call Δ a *θ-oriented* triangle at p. See Figure 1.

Let Ψ be the set of orientations $\{\frac{\alpha}{4}, \frac{2\alpha}{4}, \ldots, \frac{\alpha\lceil 8\pi/\alpha\rceil}{4}\}$. We call a triangle Δ a *θ-critical* triangle at p if Δ is a good triangle at p, and Δ is $(\theta-\frac{\alpha}{4})$-oriented at p, θ-oriented at p, and $(\theta+\frac{\alpha}{4})$-oriented at p. Observe that for each $c \in \mathcal{C}$ and $p \in \partial c$, there exist a θ-critical triangle at p, for some $\theta \in \Psi$. For each $c \in \mathcal{C}$ and each $\theta \in \Psi$ let $\gamma_\theta(c)$ denote the portion of ∂c consisting of points p such that p is $(\theta-\frac{\alpha}{4})$-oriented, θ-oriented, and ,$(\theta+\frac{\alpha}{4})$-oriented. By the constant description complexity assumption made in the introduction, $\gamma_\theta(c)$ consists of at most s connected portions of ∂c. We further divide these portions of $\gamma_\theta(c)$ into a constant number of 'not-too-long' subarcs (that might overlap), called *primitive arcs* or *p-arcs* for short. Each p-arc δ is required (i) to be differentiable (that is, there exists a well defined tangent at each relatively interior point of δ), (ii) not to contain in its relative interior any locally x-extremal or y-extremal point or any inflection point of ∂a, and (iii) to satisfy the property that the difference in the orientations of the tangents at any pair of points of δ is at most π/t, for some predetermined integer $t > 10$. A p-arc along the boundary of an object c is *convex* if the segment connecting the endpoints of the arc is contained in c. Otherwise, we say that the p-arc is *concave.*

For each $c \in C$ and for every $\theta \in \Psi$, we place a θ-oriented triangle at each endpoint of every p-arc of $\gamma_\theta(c)$, and we let P_c denote the collection of these triangles. The p-arcs are chosen sufficiently short, so that the boundary of each connected component of $c \setminus \bigcup P_c$ contains at most a single p-arc. We call a maximally connected component of $c \setminus \bigcup P_c$ a *cap*, and the segment connecting the endpoints of its p-arc the *chord* of the cap. The union of a cap and the two triangles of P_c adjacent to the endpoint of its p-arc is called a *sub-object*, see Figure 2. If a sub-object is not simply connected, we 'fill in' its holes and add them to the sub-object. The boundary of a sub-object consists of a single p-arc and of portions of edges of the good triangles of P_c adjacent to the p-arc's endpoints. Note that the chord of the sub-object is generally not part of the sub-object. The collection of all sub-objects of c that are adjacent to p-arcs that are θ-oriented is denoted by c^θ. See Figure 3. Clearly c^θ consists of a constant number of sub-objects. Let $\mathcal{C}^\theta$ denote the collection of all sub-objects with this property of every $c \in \mathcal{C}$.

Fix $\theta \in \Psi$, which we assume, for simplicity, to be the negative vertical direction, otherwise rotate the plane. Define a segment tree $\mathcal{T}_\theta$ over orthogonal the y-projections of the sub-objects of $\mathcal{C}^\theta$. Each node $\mu \in \mathcal{T}$ is associated with a subset $S_\mu \subseteq \mathcal{C}^\theta$ and with a horizontal slab I_μ.

Fix $\theta_A, \theta_B \in \Psi$, (not necessarily distinct) and levels i_A of $\mathcal{T}_{\theta_A}$ and i_B of $\mathcal{T}_{\theta_B}$. Note that there are $O(\log^2 n)$ quadruples $(\theta_A, \theta_B, i_A, i_B)$ of this kind. Define A (resp. B) to be the collection of sub-objects in S_μ for μ in the i_A'th level of $\mathcal{T}_{\theta_A}$ (resp. the i_B'th level of $\mathcal{T}_{\theta_B}$). Let $U(\theta_A, \theta_B, i_A, i_B)$ denote the set of 'mixed' vertices of $\partial \bigcup(A \cup B)$ that lie on $\gamma_{\theta_A}(a)$ for some $a \in A$, and on $\gamma_{\theta_B}(b)$ for some $b \in B$. The following section is dedicated to the proof of the following lemma.

Lemma 2.1 *The size of* $U(\theta_A, \theta_B, i_A, i_B)$ *is* $O(\lambda_{s+2}(n) \log\log n)$.

It is easy to see that the proof of Theorem 1.1 follows immediately from Lemma 2.1, because for each vertex $v \in \partial \bigcup \mathcal{C}$ there exist θ_A, $\theta_B \in \Psi$ and levels i_A, i_B such that v appears as a vertex in the corresponding set $U(\theta_A, \theta_B, i_A, i_B)$, and because the number of quadruples $(\theta_A, \theta_B, i_A, i_B)$ is $O(\log^2 n)$.

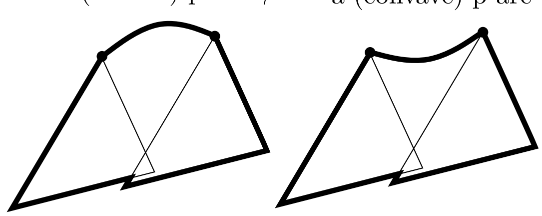

Figure 2: Two examples of sub-objects.

3 Proof of Lemma 2.1

We fix a quadruple $(\theta_A, \theta_B, i_A, i_B)$, as above. We verify that $\theta_A = 3\pi/2$ (the negative y-direction) by rotating the plane if necessary. Let μ be a node in the i_A'th level of $\mathcal{T}_{3\pi/2}$, and let I_μ be the horizontal slab associated with μ. Let c be a subobject of S_μ, let $p \in \gamma_{3\pi/2}(c) \cap I_\mu$ and let Δ be a $(3\pi/2)$-oriented triangle of c at p. Note that Δ is an α-fat triangle, and the length of the y-span of c is at least the width of I_μ. Hence there is a constant integer l,

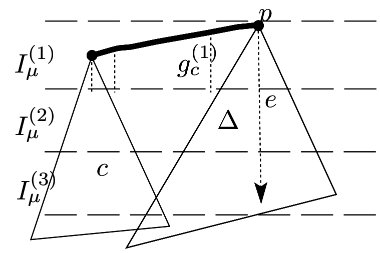

Figure 3: The slab I_μ, the strip to which I_μ is split, and the function $g_c^{(i)}(x)$

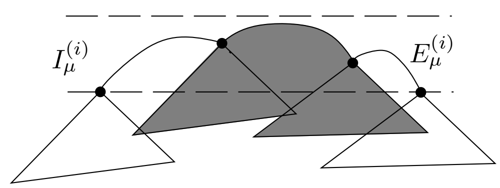

Figure 4: The upper envelope $E_\mu^{(i)}$ and some new sub-objects defined by its vertices. One of the subobjects is shaded.

such that the length of the y-span of each edge of Δ is at least $1/l$ times the width of I_μ. We divide I_μ into a constant number of *strips*, $I_\mu^{(1)}, \ldots, I_\mu^{(l)}$, of equal width. Thus, p lies in a different strip than the other two vertices of Δ, See Figure 3. Thus, for each strip $I_\mu^{(i)}$, we can express all points of $I_\mu^{(i)} \cap \gamma_{3\pi/2}(c)$ as a graph of a function $g_c^{(i)}(x)$ defined on the lower boundary of the strip $I_\mu^{(i)}$.

For each strip $I_\mu^{(i)}$, consider the upper envelope $E_\mu^{(i)}$ (See Figure 4)) of the functions $g_c^{(i)}(x)$, for $c \in \mathcal{S}_\mu$. Let E_μ denote the union of these upper envelopes for all strips of I_μ, and let E_A denote the union of all these envelopes, taken over all nodes μ in the i_A'th level of $\mathcal{T}_{\theta_A}$. Repeat the same analysis for θ_B, and obtain a corresponding union E_B of upper envelopes (relative to the θ_B-direction).

Let v be a vertex of E_A incident to the boundaries of sub-objects $c_1, c_2 \in A$. We add to P_{c_1} a θ_A-critical triangle for c_1 at v, and to P_{c_2} a θ_A-critical triangle for c_2 at v (with an appropriate construction, these are similar triangles with a common vertex and overlapping edges). For each $c \in \mathcal{S}_\mu$ and each p-arc γ of c we add θ_A-critical triangles at each point where γ crosses a boundary of a strip of I_μ. We further refine the splitting of arcs into p-arcs, so that no p-arc γ contains a vertex of any of the new triangles, except of course for its endpoints. Sub-objects are split as well, so that each sub-object contains exactly one new p-arc on its boundary. Observe that now each p-arc is contained in at most one strip of I_μ.

We next remove from A all sub-objects that do not participate in E_A. Thus each p-arc of a remaining sub-object of A is fully contained in E_A, and also fully contained in a single strip of some I_μ. Analogously, we restructure the sub-objects and p-arcs for B, the collections $\{P_b\}_{b\in B}$ and the union of envelopes E_B. We list several important attributes of this construction:

(A1) Two p-arcs of A (resp. B) are either disjoint, or intersect only at their endpoints. Moreover, a p-arc γ of a sub-object $a_1 \in A$ might intersect the boundary of a different sub-object $a_2 \in A$ only at an endpoint of γ, or at a point of ∂P_{a_2}. Similar attributes hold for B.

(A2) A necessary condition for a vertex v to belong to $U(\theta_A, \theta_B, i_A, i_B)$ is that v lies on E_A and on E_B.

(A3) The complexity of E_A and of E_B are each $O(\lambda_{s+2}(n))$.

Consider the collections $P_A = \bigcup_{a\in A} P_a$, $P_B = \bigcup_{b\in B} P_b$. The result of [12] implies that the complexity of $\partial \bigcup P_A$ and of $\partial \bigcup P_B$ are each $O(\lambda_{s+2}(n)\log\log n)$, as each triangle in these collections is an α-fat triangle (the constants of proportionality depend on α). Define $UP(A)$ as the set of all vertices that are either vertices of $E_A \cap \partial \bigcup(A \cup P_A \cup P_B)$, or vertices of sub-objects of A or vertices of triangles in $P_A \cup P_B$. We define $UP(B)$ in a fully symmetric manner, interchanging A and B.

We first state a slightly modified version of a lemma that appeared in [4]. The proof is deleted from this extended abstract.

Lemma 3.1 *[Efrat & Sharir, 97] Let K_a be the portion of a cap of some sub-object $a \in \mathcal{C}^\theta$, enclosed between its p-arc γ_a and its chord e_a, such that γ_a is convex. Let $\Delta_b \in P_b$ be a good triangle for some object $b \in \mathcal{C}$, such that the edge e_b of Δ_b crosses γ_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 that is an endpoint of e_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 $K_a \cap \Delta_b$.

Lemma 3.2 *The number of vertices of $UP(A)$ and of $UP(B)$ is $O(\lambda_{s+2}(n)\log\log n)$.*

Proof: It suffices to prove the lemma for $UP(B)$. Let v be a vertex of $UP(B)$, lying on an edge e of a triangle Δ in $P_A \cup P_B$ and on a p-arc $\gamma = \gamma_b$ contained in $\gamma_{\theta_B}(b)$, for some $b \in B$. (All other kinds of vertices are trivial to bound.) Let u_1, u_2 be the endpoints of γ_b; See Figure 6. Assume again that $\theta_B = 3\pi/2$, so the slabs of $\mathcal{T}_{\theta_B}$ are horizontal. Let μ be the node of $\mathcal{T}_{\theta_B}$, in the i_B'th level, associated with the sub-object containing v on its boundary. Let t_1 and t_2 be the triangles of P_b, which are θ_B-oriented for b at u_1 and at u_2. Let F be the axis-parallel rectangle formed by intersecting $I_\mu^{(i)}$ with the vertical strip spanned by γ_b (see Figure 6). Clearly v lies in F. If $\Delta \cap \gamma_b$ fully contains one of the two portions of γ_b connecting v to one of its endpoints, we charge v to this endpoint. Since the

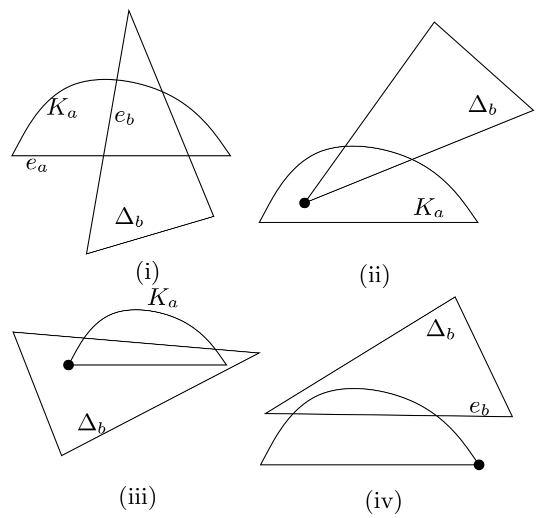

Figure 5: Illustrating the various cases in Lemma 3.1

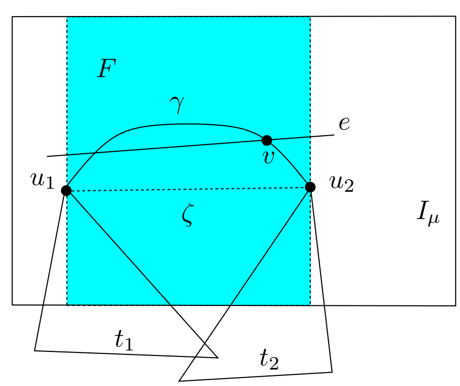

Figure 6: Illustrating the proof of Lemma 3.2; e is long

number of endpoints is $O(\lambda_{s+2}(n))$ and each can be charged at most twice, the number of vertices v of this kind is within the asserted bound. So assume this is not the case.

Recall that γ_b is either concave or convex, so e intersects γ_b either once or twice. We call e a *long* edge if both its endpoints are outside F; otherwise e is a *short* edge. If e is short, we uniquely charge v to one of the endpoints of e inside F; again, the number of such endpoints is within the asserted bound. So let us assume that e is long.

If e intersects γ_b once, then either there is an endpoint of e inside the cap of γ_b, or (since e is long) it must intersect t_1, t_2, or some other triangle of $P_A \cup P_B$, at a point inside the cap and on $\partial \bigcup (P_A \cup P_B)$, so we can charge v to this intersection point (and the number of such intersections is within the asserted bound). So assume that e intersects γ_b twice. If γ_b is concave, then if we trace e from v into b, we reach a vertex of $\bigcup (P_A \cup P_B)$, to which we can charge v. So we may assume that γ_b is convex, as depicted in Figure 6.

Let Δ be the triangle of P_a incident to e, and let z be the vertex of Δ that lies opposite to e. We say that e is *special* if z lies inside F. Since we can charge v in this case to z, it suffices to consider the case where e is non-special and long.

Applying Lemma 3.2 to e and the appropriate cap portion, we see that if any of the cases (i)–(iii) arises, we can charge v to a vertex of $\partial \bigcup (P_A \cup P_B)$ inside the cap, as done above. So we may assume that case (iv) arises.

We now claim that the number of long non-special edges $e_1, \ldots, e_l$ incident to vertices on γ_b and satisfying property (iv) of Lemma 3.2 is a constant. Indeed, let $\Gamma(e_i)$ be the portion of γ_b spanned between its two intersection points with e_i. It is impossible that $\Gamma(e_i)$ and

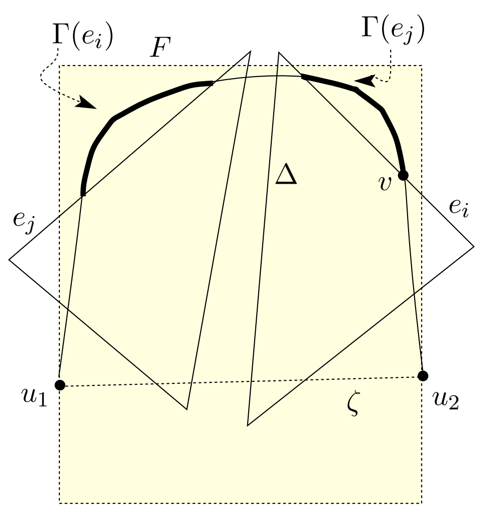

Figure 7: If $\Gamma(e_i)$ and $\Gamma(e_j)$ are disjoint then e_i and e_j cross different pairs of edges of F.

$\Gamma(e_j)$ intersect. Indeed if $\Gamma(e_i) \cap \Gamma(e_j) \neq \emptyset$, but neither $\Gamma(e_i) \subseteq \Gamma(e_j)$ nor $\Gamma(e_j) \subseteq \Gamma(e_j)$, then e_i and e_j must intersect inside F (by the convexity of γ), thus they are not long. On the other hand, it is impossible that one of them fully contains the other since they both satisfy property (iv). Moreover, if $\Gamma(e_i)$ and $\Gamma(e_j)$ are disjoint (see Figure 7), then it is easily verified that e_i and e_j intersect different pairs of edges of F. This concludes the proof of the lemma. □

We can now turn to the proof of Lemma 2.1. A vertex v of $\partial(a \cap b)$ is an *irregular vertex* if the number of vertices in the connected component of $a \cap b$ incident to v is at least 4. Otherwise, v is a *regular vertex*. We refer the reader to a slightly different definition of regular and irregular vertices, and relevant combinatorial results, in [1] and [13].

We first bound the number of irregular vertices of $U(\theta_A, \theta_B, i_A, i_B)$.

Lemma 3.3 *The number of irregular vertices of* $U(\theta_A, \theta_B, i_A, i_B)$ *is* $O(\lambda_{s+2}(n) \log\log n)$.

Proof: Let v be an irregular vertex of $U(\theta_A, \theta_B, i_A, i_B)$, incident to a p-arc γ_a and to another p-arc γ_b, for some $a \in A$, $b \in B$, Let I_μ (for $\mu \in \mathcal{T}_{\theta_B}$) be the strip containing b and assume again that $\theta_B = 3\pi/2$, otherwise rotate the plane. As in Lemma 3.2, let F be the rectangle formed by the intersection of I_μ with the vertical strip spanned by γ_b. We call γ_a *special* if $\gamma_a \cap F$ contains either an endpoint of γ_a, or a locally highest point, or a locally lowest point or a locally rightmost point, or a locally leftmost point of γ_a. Note that if γ_a is special, we can charge v to one of the extreme points listed above, since there is only a constant number of them on each object of $\mathcal{C}$.

Let μ_a and μ_b be the normals at v to γ_a and γ_b, pointing into a and b, respectively. Let ϕ be the smaller angle between μ_a and μ_b. Let $\phi_0 < \pi/10$ denote the maximal turning angle of any p-arc.

We distinguish between three cases:

* $\phi_0 \leq \phi < \pi - \phi_0$ (see Figure 8(i)). Clearly, in this case γ_a and γ_b have at most one intersection point, which must be v itself. Indeed, construct a line ℓ that passes through v, and forms angles $\phi/2$ and $-\phi/2$ with γ_a and γ_b, respectively. Since neither γ_a nor γ_b can turn by more than ϕ, it follows that, apart from v, ℓ is disjoint from both γ_a and γ_b, so, apart from v, ℓ separates γ_a from γ_b.

We follow γ_a from v in the direction in which it enters b. Since γ_a has entered the cap of b bounded by γ_b and it does not intersect γ_b again, it either ends within the cap or meets a triangle in $P_A \cup P_B$. In either case we can charge v to this endpoint or intersection point. (Note that in the latter case, this intersection must be a vertex of $UP(A)$.)

* $\phi > \pi - \phi_0$ (see Figure 8(ii)). Without loss of generality, assume that the situation is as shown in Figure 8(ii). That is, a lies above γ_a near v and b lies below γ_b near v, and as we trace γ_a and γ_b to the left, each of them enters into the other object. If we reach in any of these tracings a point on $\bigcup(P_A \cup P_B)$ then this is a vertex of either $UP(A)$ or $UP(B)$, to which we can charge v. So assume this is not the case. Hence, γ_a and γ_b must intersect again. It is obvious from the condition on the angles and the assumptions made so far that in this case v is a regular vertex. We deal with this type of vertices later on.

* $\phi < \phi_0$ (see Figure 8(iii)). This case is more involved. Observe that the tangent of every point of γ_b is "almost horizontal", that is, its orientation is in the range $(-\phi_0, \phi_0)$. Thus the orientation of every point of γ_α is inside F is in the range $(-2\phi_0, 2\phi_0)$. Thus the triangles of a are $(3\pi/2)$-oriented as well. Let q_l, q_r be the left and right endpoints of γ_b.

We follow γ_a from v in the direction inward b — say to the left (see Figure 9). If we

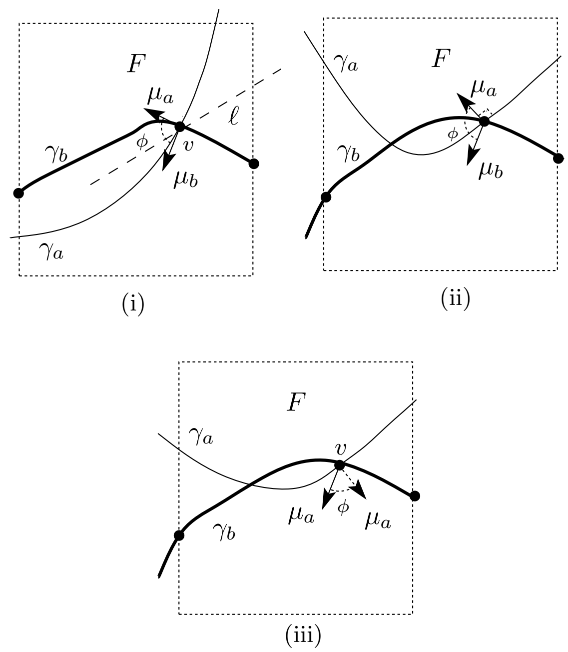

Figure 8:

reach a triangle of $P_A \cup P_B$, then this is a vertex u of $UP(A)$ that we charge. Thus we may assume that we reach another vertex v_2 on $\gamma_b \cap \gamma_a$, to the left of v. By attribute A1 we are guaranteed that we have not enter an sub-object $a' \in A$ so far (in the portion of γ_a between v and v_2). Hence we deduce that v_2 is also a point of $U(\theta_A, \theta_B, i_A, i_B)$.

Let J_r (resp. J_l) be the portion of $\gamma_b \cap a$ adjacent to v (resp. v_2). If J_l or J_r contain a point of $P_A \cup P_B$, then this is a vertex u of $UP(B)$, which we can charge to v and to v_2, as u can be charged only a constant number of times in this way. Assuming this is not the case. Let L_l and L_r be the portions of $\gamma_a \setminus b$ adjacent to v_r and v respectively. Assume that $L_l, (resp.\ L_r)$ passes above q_l, the left (resp. q_r, the right) endpoint of γ_b. It is not hard to verify, by the way p-arcs were defined and the fact that the triangles of a are $(3\pi/2)$-oriented (though not necessarily $(3\pi/2)$-critical), that either q_l (resp. q_r) is inside a, or that by tracing γ_b to left (resp. right) from v_2 (resp. v), we must encounter one of the triangles of P_a, in a vertex of $UP(B)$. In the former case, we deduce that v_2 (resp, v) is the leftmost (resp. rightmost) vertex of $U(\theta_A, \theta_B, i_A, i_B)$ on γ_b, and we charge both v and v_2 to this endpoint in this case. Thus we assume that this is not the case.

If $L_l \cap F$ or $L_r \cap F$ contain a point of $P_A \cup P_B$, then as above, this a vertex of $UP(A)$ that we can charge, as it is inside F. If on the other hand both L_a and L_b intersect the roof edge of ∂F, then the portion of γ_a between v and v_2 must contain a locally minimal point u (lowest point) which is inside b, and we can charge u to v and v_2. (In the case that u lies inside F, then this is also a contradiction to our assumption that γ_a is not special.) Hence at least one of L_l and L_r, say L_l, lies completely inside F. If L_r but not L_l lies completely inside F, we reverse the direction by which we traverse γ_a. The other endpoint v_3 of L_l (or L_r if we have reversed the direction) must therefor also be a vertex of $U(\theta_A, \theta_B, i_A, i_B)$, otherwise γ_a would have a point of $P_A \cup P_B$ inside F. We continue following γ_a along γ_a to the left direction, possibly meeting more vertices of $U(\theta_A, \theta_B, i_A, i_B)$ that belong to $\gamma_a \cap \gamma_b$. Their number however is $\leq s$. Thus this process must end after discovering at most s vertices, and since the only way that the process ends is that we discover a vertex that we can charge (to all $\leq s$ vertices of $\gamma_a \cap \gamma_b$), we have obtained a bound on the number of irregular vertices of $U(\theta_A, \theta_B, i_A, i_B)$ in this case as well. This concludes the proof of Lemma 3.3. □

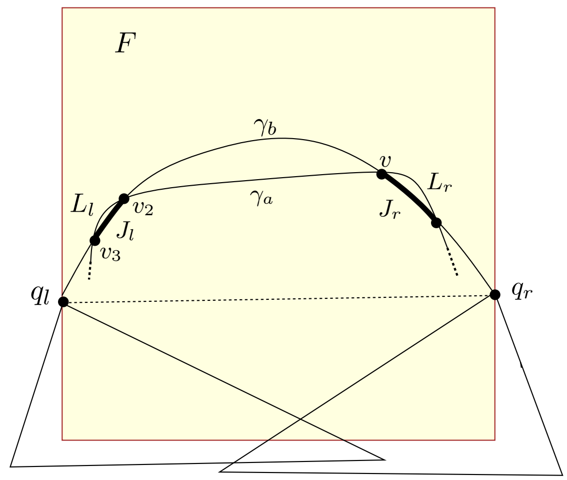

Figure 9: The third case of the proof of Lemma 3.3.

To complete the proof of Lemma 2.1, it remains to bound the number of regular vertices in $U(\theta_A, \theta_B, i_A, i_B)$. Set $b \in B$, and assume that $\gamma_1, \ldots, \gamma_l$ are all the non-special p-arcs belonging to (not necessarily distinct) respective sub-objects $a_1, \ldots, a_l$ of A, each containing a regular vertex v_i of $U(\theta_A, \theta_B, i_A, i_B)$ that lies on γ_b. For each i, define $\Gamma(\gamma_i)$ as the portion of $\gamma_b \cap a_i$ incident to v_i. Note that γ_i cannot intersect γ_j inside F (for any $1 \leq i < j \leq l$) because each γ_i is a non-special p-arc. It is not hard to show that there can be only a constant number of pairs of p-arcs γ_i, γ_j, such that $\Gamma(\gamma_i) \cap \Gamma(\gamma_j)$ is empty, since each such pair must intersect different edges of ∂F. Similarly there is no pair γ_i, γ_i such that $\Gamma(\gamma_i) \cap \Gamma(\gamma_j)$ partially overlap (that is, $\Gamma(\gamma_i) \cap \Gamma(\gamma_j) \neq \emptyset$ but neither $\Gamma(\gamma_i) \subseteq \Gamma(\gamma_j)$ nor $\Gamma(\gamma_j) \subseteq \Gamma(\gamma_i)$). On the other hand, it is impossible that $\Gamma(\gamma_i) \subseteq \Gamma(\gamma_j)$, since this would imply that v_j is not a regular vertex. Indeed, the p-arc γ_b passes through γ_i to create v_i, gets out of a_i in order to meet a_j at v_j, and returns to a_i, which implies that there are at least 4 intersection points in the same connected component of $a \cap b$. This contradiction concludes the proof of Lemma 2.1. □

Remark 3.4: The bound of Theorem 1.1 improves if objects of $\mathcal{C}$ are roughly the same size. Assume that there are constants d, κ, such that $d \leq diam(c) \leq \kappa d$, for each object of $c \in \mathcal{C}$. Then the bounds of Theorem 1.1 improve to $O(\lambda_{s+2}(n))$. This follows by modifying the preceding proof, and we will only comment on a few of the less trivial modifications that are required.

For each orientation $\theta_A \in \Psi$, we divide the plane into infinite parallel strips of width

μd (for a sufficiently small constant μ that depends on s, α and β), orthogonal to the θ_A direction, such that (as above) if Δ is a θ_A-critical triangle to an object c at a point $p \in \partial c$, then the other two vertices of Δ do not lie in the strip containing p. We define A as the union of sub-objects incident to p-arcs of $\gamma_{\theta_A}(c)$, over all $c \in \mathcal{C}$. The definitions of B and of all the other notations used in the proof are analogous. We also use the fact that all the oriented triangles in P_A and P_B are roughly of the same size, and thus the complexity of their union is only $O(n)$, as shown in [2].

4 Conclusions remarks

The definition of (α, β)-covered object is not the first attempt to define fatness for non-convex objects; In [16], Frank van der Stappen gives the following definition to fatness. An object $C \subseteq \mathbb{R}^d$ is δ-fat (for $0 < \delta < 1$) if for each d-dimensional ball B, whose center is inside C but does not contain C completely, that the volume of $(B \cap C)$ is at least the volume of B. Two questions that naturally arises, are (i) what is the relation between (α, β)-covered object and δ-fat objects, and (ii), can one shows a bound on the complexity of the union of δ-fat objects. Recently van der Stappen [15] answered both these questions: He showed that the definition of δ-fat object is stronger than the definition of (α, β)-objects by showing that each (α, β)-covered object is also a δ-fat object, for an appropriate parameter δ (that depends on α and β). He also answered the second question by presenting a construction showing that the boundary of the union of n δ-fat objects can has $\Omega(n^2)$ vertices, implying that δ-fatness is not suffices to provided sub-quadratic complexity.

Acknowledgment.

We wish to thank Micha Sharir for helpful discussions and many useful ideas regards the results of this paper. Thanks also to anonymous referee for many useful remarks.

References

[1] B. Aronov, A. Efrat D. Halperin and M. Sharir, A subquadratic bound on the number of regular vertices of the union of Jordan regions. *Proceedings 6$^{\text{th}}$Scand. Workshop on Algorithms Theory* (1998), 322–334.

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

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

[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 and M. Katz, On the union of α-curved objects, *Proceedings 14$^{\text{th}}$Annual Symposium on Computational Geometry*, 1998, 206–213.

[6] A. Efrat, M.J. Katz, F. Nielsen and M. Sharir, Dynamic data structures for fat objects and their applications. *Proc. 5rd Workshop on Algorithms and Data Structures*, (1997), 297–396.

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

[8] M. Katz, 3-D vertical ray shooting and 2-D point enclosure, range searching, and arc shooting amidst convex fat objects, *Computational Geometry: Theory and Applications* 8 (1997), 299–316.

[9] 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.

[10] 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.

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

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

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

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

[15] A.F. van der Stappen, private communication, June 1999.

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