Discrete Comput Geom (2008) 40: 127–140
DOI 10.1007/s00454-007-9029-7

Improved Bounds on the Union Complexity of Fat Objects

Mark de Berg

Received: 20 October 2006 / Revised: 9 July 2007 /
Published online: 15 September 2007

Abstract We introduce a new class of fat, not necessarily convex or polygonal, objects in the plane, namely locally γ-fat objects. We prove that the union complexity of any set of n such objects is $O(\lambda_{s+2}(n)\log^2 n)$. This improves the best known bound, and extends it to a more general class of objects.

Keywords Fat Objects · Union Complexity · Density Lemma

1 Introduction

The running time of geometric algorithms and the amount of storage used by geometric data structures often depend on the combinatorial complexity of certain geometric structures. Hence, the study of the combinatorial complexity of geometric structures is an important and active area within computational geometry. In this paper we study the combinatorial complexity of the union of a set $\mathcal{F}$ of n objects in the plane. This is relevant because there are many geometric algorithms and data structures whose performance depends on the union complexity of certain collections of planar objects. Examples are algorithms for hidden-surface removal [13], data structures for ray shooting [4, 5, 12], algorithms for computing depth orders [1, 5, 12], and algorithms for motion planning [14, 20].

In the worst case the complexity of the union of n constant-complexity objects in the plane can be as high as $\Theta(n^2)$, a bound which is for example achieved by a set of n long and thin rectangles arranged in a grid-like pattern. In many applications,

The research was supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 639.023.301.

M. de Berg (✉)
Department of Computing Science, TU Eindhoven, P.O. Box 513, 5600 MB Eindhoven, The Netherlands
e-mail: mdberg@win.tue.nl

however, one would expect that the objects have some favorable properties and that the union complexity is much lower. One such property that has received considerable attention is *fatness*. Intuitively, an object is called fat if it is not arbitrarily long and skinny—see Sect. 2 for precise definitions. There are many algorithmic results for fat objects, several of which depend on the union complexity of fat objects in the plane. Hence, the union complexity of fat objects in the plane has been studied extensively.

One of the first results on the union complexity of fat objects was for fat wedges, that is, wedges whose interior angle is bounded from below by a constant. For this case it has been shown [2, 10] that the union complexity is $O(n)$. Matoušek et al. [15] considered the case of *δ-fat triangles*, that is, triangles all of whose angles are at least δ for some fixed constant δ. They proved that the union complexity of n such triangles is $O((1/\delta^3)n \log\log n)$. Later this bound was improved by Pach and Tardos [17] to $O((1/\delta \log(1/\delta))n \log\log n)$.

Several people have worked on extending these results to more general types of fat objects, in particular to curved and/or non-convex objects [8, 9, 11, 21] and to higher dimensions [3, 16]. The most general result for planar objects to date is by Efrat [8], who considered so-called (α, β)-covered objects—see the next section for a definition. Efrat proved that the union complexity of n constant-complexity (α, β)-covered objects is bounded by $O(\lambda_{s+2}(n) \log^2 n \log\log n)$, where s is the maximum number of intersections between any pair of object boundaries and $\lambda_t(n)$ denotes the maximum length of an (n, t) Davenport–Schinzel sequence; $\lambda_t(n)$ is near-linear for any constant t [18]. (The constant hidden in the O-notation depends on the constants α and β.)

We introduce in Sect. 2 a new class of fat objects in the plane, namely *locally γ-fat objects*. This class is more general than the class of (α, β)-covered objects. We prove that the union complexity of n constant-complexity locally γ-fat objects is $O(\lambda_{s+2}(n) \log^2 n)$, thus not only generalizing the result of Efrat but also slightly improving the bound.

One of the novel ingredients of our proof is a lemma—we call it the Density Lemma—stating that the set of boundary pieces of the union of fat objects in the plane has low density. This is a powerful tool, which we also used elsewhere [4], and which we believe will find other applications in the study of unions of fat objects in the plane. The Density Lemma implies that the union complexity of two sets of locally fat objects can be bounded by the sum of the union complexities of the two sets (with a constant multiplicative factor that depends on the fatness). It then follows for example that the union of a set of fat triangles and disks in the plane has $O(n \log\log n)$ complexity, which was, to the best of our knowledge, not known.

The Density Lemma allows us to reduce the problem of bounding the union complexity of locally γ-fat objects to the problem of bounding the union complexity of so-called consistently oriented fat quasi-triangles; these are almost triangular fat shapes with two edges in a fixed orientation. We then give a simple proof that the union complexity of such shapes is $O(\lambda_{s+2}(n) \log^2 n)$. An interesting feature of our proof is that, unlike Efrat's proof, it does not rely on the result of Matoušek et al. [15] for fat triangles.

2 Preliminaries

Let $\mathcal{F} := \{o_1, \ldots, o_n\}$ be a set of objects in the plane. From now on, we assume that each object is compact, that is, bounded and closed. We also assume that each object has constant complexity; in particular we assume that the boundary of each object consists of $O(1)$ algebraic curves of constant maximum degree. Hence, any two object boundaries intersect at most s times for some constant s. We denote the boundary of an object o by ∂o.

Fatness and low density We first define (α, β)-covered objects, as introduced by Efrat [8]. Figure 1(i) illustrates the definition.

Definition 2.1 A planar object o is called (α, β)*-covered* if for every point p on the boundary of o one can place a triangle t_p with the following properties:

(i) t_p is contained inside o;
(ii) p is a vertex of t_p;
(iii) t_p is an α-fat triangle, that is, all its angles are at least α;
(iv) the length of each edge of t_p is at least $\beta \cdot diam(o)$, where $diam(o)$ is the diameter of o.

Next we introduce a new characterization of fatness. Let $area(o)$ denote the area of an object o. Consider an object o and a disk D whose center lies inside o. If o is non-convex $D \cap o$ may consist of several connected components. We define $D \sqcap o$ to be the connected component that contains the center of D.

Definition 2.2 Let o be an object in the plane and γ a parameter with $0 \leq \gamma \leq 1$. We say that o is *locally* γ*-fat* if, for any disk D whose center lies in o and that does not fully contain o in its interior, we have $area(D \sqcap o) \geq \gamma \cdot area(D)$.

This definition is illustrated in Fig. 1(ii). It is similar to the fatness definition introduced by Van der Stappen et al. [19, 20], except that we use $area(D \sqcap o)$ instead of $area(D \cap o)$. Thus for convex objects, where $D \sqcap o = D \cap o$, the definitions are identical.

The following proposition shows the relation between (α, β)-covered objects and locally γ-fat objects.

Proposition 2.3

(i) *Any* (α, β)*-covered object is locally* γ*-fat for some* $\gamma = \Omega(\alpha\beta^2)$.
(ii) *There are no constants* α, β *that depend only on* γ *such that any locally* γ*-fat object is an* (α, β)*-covered object.*

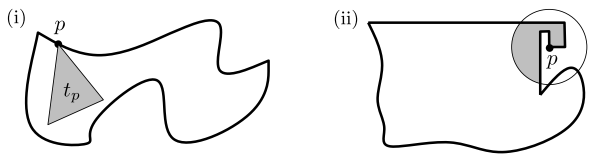

Fig. 1 Illustration of the definition of (α, β)-covered object and the definition of locally γ-fat object

Fig. 2 (**i**) A fat object that is not locally fat. (**ii**) The first two objects in a construction showing that objects that are fat according to the definition of Van der Stappen et al. [19, 20] can have a quadratic union

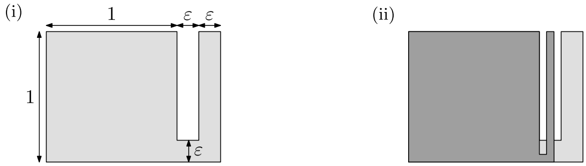

Proof (i) Let o be an (α,β)-covered object and let D be a disk centered at a point $p \in o$ and not containing o in its interior.

First assume $p \in \partial o$. Then there is an α-fat triangle $t_p \subset o$ with p as a vertex all of whose edges have length at least $\beta \cdot diam(o)$. Clearly $(D \cap t_p) \subset (D \sqcap o)$. If t_p is not fully contained in D then $area(D \cap t_p) = \Omega(\alpha \cdot area(D))$ because all angles of t_p are at least α. If, on the other hand, $t_p \subset D$ then

$$area(D \sqcap o) \geq area(D \cap t_p) = area(t_p) = \Omega\big(\alpha \cdot (\beta \cdot diam(o))^2\big) = \Omega\big(\alpha\beta^2 \cdot area(D)\big),$$

where the last equality follows because D does not fully contain o.

Now assume p lies in the interior of o. Let p' be a point on ∂o with minimum distance to p. If $dist(p, p') \geq radius(D)/2$ then $area(D \sqcap o) \geq area(D)/4$. Otherwise, let $D' \subset D$ be the disk with center p' and radius $radius(D)/2$. Now $(D' \sqcap o) \subset (D \sqcap o)$, and since $p' \in \partial o$ we can apply the argument above and get

$$area(D \sqcap o) \geq area(D' \sqcap o) = \Omega\big(\alpha\beta^2 \cdot area(D')\big) = \Omega\big(\alpha\beta^2 \cdot area(D)\big).$$

(ii) Consider the object shown in Fig. 1(ii). It is impossible to place a triangle with the point p as a vertex that is relatively large and stays inside the object. In fact, the size of the largest triangle that can be placed at p will tend to zero as the part of the object sticking out in the top right shrinks, while the local fatness does not significantly change when this happens. □

Thus the class of locally γ-fat objects is more general than the class of (α,β)-covered objects. On the other hand, the class of locally fat objects is less general than the class of fat objects as defined by Van der Stappen et al. The object in Fig. 2(i) shows this: when ε tends to zero, the local fatness of the object tends to zero as well, while the fatness according to the definition of Van der Stappen et al. remains lower bounded by a fixed constant. A more strict definition is a necessity, however, to obtain near-linear bounds on the union complexity: one can place n objects of the type shown in Fig. 2(i) in such a way that the resulting union has complexity $\Omega(n^2)$. To see this, consider Fig. 2(ii). It consists of two such objects, where the value of ε in the darker object is three times as small as in the lighter object. We can put another $(n/2) - 2$ objects in a similar way, each time reducing the value of ε by a factor three, to obtain a collection of objects whose union is a square with $n/2$ thin vertical rectangles to the right of it. We can then copy this whole construction and rotate it by 90 degrees to obtain a union of complexity $\Omega(n^2)$.

Besides the concept of fatness, we also need the concept of *density* [6]. For an object o in $\mathbb{R}^2$, we use $size(o)$ to denote the radius of the smallest enclosing disk of o. Note that a locally γ-fat object o in $\mathbb{R}^2$ has area $\Omega(\gamma \cdot size(o)^2)$.

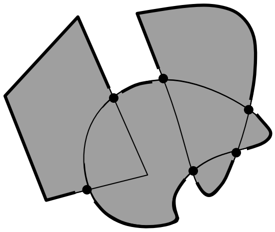

Fig. 3 A collection of three objects whose union has complexity 12

Definition 2.4 The *density* of a set S of objects in $\mathbb{R}^2$ is defined as the smallest number λ such that the following holds: any disk $D \subset \mathbb{R}^2$ is intersected by at most λ objects $o \in S$ such that $size(o) \geq size(D)$.

Notation We end with some more notation and terminology. We use $\mathcal{U}(S)$ to denote the union of a set S of objects. The union boundary $\partial\mathcal{U}(S)$ of a set S of planar objects consists of maximally connected portions of the boundaries of the objects in S. We call these portions the *edges* of $\mathcal{U}(S)$; the endpoints of these edges are called the *corners* of $\mathcal{U}(S)$.[1] The *(combinatorial) complexity* of $\mathcal{U}(S)$ is defined as the total number of edges and corners of $\mathcal{U}(S)$. For example, the union in Fig. 3 has complexity 12, as it has six corners and six edges. Notice that, up to an additive term equal to the number of objects in S, the complexity of $\mathcal{U}(S)$ is linear in the number of corners. Hence, it suffices to bound that number.

3 The Density Lemma

In this section we prove the Density Lemma, which will enable us to bound the complexity of the union of two unions of fat objects. Recall that by an edge of the union of a set of objects we mean a maximally connected portion of the union boundary that is contributed by a single object.

Lemma 3.1 (Density Lemma) *Let $\mathcal{F}$ be a set of n locally γ-fat objects, and let $E(\mathcal{F})$ denote the set of edges of the union $\mathcal{U}(\mathcal{F})$. Then the density of $E(\mathcal{F})$ is $O(1/\gamma)$.*

Proof Let D be a disk, and assume without loss of generality that $size(D) = 1$. Let $E_D \subset E(\mathcal{F})$ be the set of edges $e \in E(\mathcal{F})$ that intersect D and for which $size(e) \geq 1$. We have to show that $|E_D| = O(1/\gamma)$.

We partition the bounding square of D into four unit squares. Together these four squares—the four squares drawn with thick lines in Fig. 4(i)—cover D. Let S_1 be any one of these squares, and let $E_{S_1} \subset E_D$ denote the edges intersecting S_1. Let S_2 and S_{mid} be squares with the same center as S_1, where S_2 has edge length $\sqrt{2}$ and S_{mid} has

[1] The boundaries of the objects may contain vertices (breakpoints between adjacent boundary segments or arcs) as well. Such vertices may also show up on the union boundary. These are not corners in our definition and, hence, do not contribute to the complexity. However, their total number is bounded by the total complexity of the objects, so counting them does not change the bounds asymptotically.

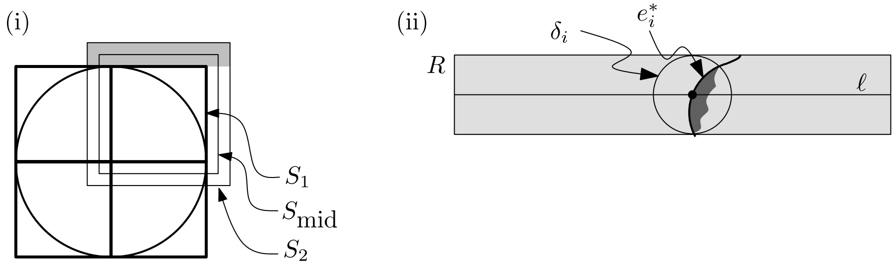

Fig. 4 Illustration for the proof of Lemma 3.1

edge length $(1+\sqrt{2})/2$. Thus S_{mid} is midway between S_1 and S_2. Since $size(e_i) \geq 1$ for any $e_i \in E_{S_1}$, such an edge e_i cannot be completely contained in the interior of S_2. Since e_i intersects S_1 by definition, it must therefore cross the square annulus $S_2 \setminus S_1$. In fact, if we cover $S_2 \setminus S_1$ using four (partially overlapping) rectangles—Fig. 4(i) shows one of these rectangles shaded—of size $\sqrt{2}$ by $(\sqrt{2}-1)/2$, then there must be one such rectangle R crossed by e_i. That is, a portion e_i^* of e_i connects the two longer sides of R—see Fig. 4(ii).

We shall bound the number of such edge portions e_i^* for which $\mathcal{U}(\mathcal{F})$ lies locally to the right of e_i^*; the number of edges for which $\mathcal{U}(\mathcal{F})$ lies locally to the left can be bounded similarly. Let p_i^* be a point where e_i^* intersects the line ℓ midway between the two longer sides of R, and let δ_i be the disk centered at p_i^* of radius $(\sqrt{2}-1)/4$. Let $object(e_i)$ be the object of which e_i is a boundary piece. Because e_i is locally γ-fat, we have

$$area(\delta_i \sqcap object(e_i)) \geq \gamma \cdot \pi \left((\sqrt{2}-1)/4\right)^2.$$

Now consider another edge $e_j \in E_{S_1}$, with a portion e_j^* crossing R and where $\mathcal{U}(\mathcal{F})$ lies locally to the right, and consider a disk δ_j of radius $(\sqrt{2}-1)/4$ centered at a point of $e_j^* \cap \ell$. Then $area(\delta_j \sqcap object(e_j))$ can be bounded as above. Because e_i^* and e_j^* are portions of union edges, they do not intersect any other object (except possibly at their endpoints). It follows that

$$(\delta_i \sqcap object(e_i)) \cap (\delta_j \sqcap object(e_j)) = \emptyset.$$

Since the area of R is $1 - \frac{1}{2}\sqrt{2}$, the number of edge portions we have to count for R for which $\mathcal{U}(\mathcal{F})$ lies locally to the right is at most

$$\frac{1-(1/2)\sqrt{2}}{\gamma \cdot \pi((\sqrt{2}-1)/4)^2} = O(1/\gamma).$$

The lemma follows. □

The Density Lemma allows us to bound the complexity of the combined union of sets of locally fat objects.

Lemma 3.2 (Merging Lemma) *Let S_1 and S_2 be sets of constant-complexity objects in the plane such that all objects in S_1 are locally γ_1-fat and all objects in S_2 are* *locally γ_2-fat. Let U_1 and U_2 denote the complexity of the union of S_1 and S_2, respectively. Then the complexity of $\mathcal{U}(S_1 \cup S_2)$ is $O(U_1/\gamma_2 + U_2/\gamma_1)$.*

Proof A corner of $\mathcal{U}(S_1 \cup S_2)$ is either a corner of $\mathcal{U}(S_1)$, or a corner of $\mathcal{U}(S_2)$, or it is the intersection of an edge of $\mathcal{U}(S_1)$ with an edge of $\mathcal{U}(S_1)$. The total number of corners of the first two types is obviously $O(U_1 + U_2)$, so it remains to count the number of corners of the third type. We charge each such corner v to the smaller of the two edges that define it, that is, if $v = e_1 \cap e_2$ and $size(e_1) \le size(e_2)$, then we charge v to e_1. Note that an edge e of $\mathcal{U}(S_1)$ is charged only by larger edges from $\mathcal{U}(S_2)$. By the Density Lemma, the set of edges of $\mathcal{U}(S_2)$ has density $O(1/\gamma_2)$. This implies that any edge from $\mathcal{U}(S_1)$ is charged $O(1/\gamma_2)$ times. A similar argument shows that any edge from $\mathcal{U}(S_2)$ is charged $O(1/\gamma_1)$ times, which proves the lemma. □

Since the union-complexity of n δ-fat triangles is bounded by $O((\frac{1}{\delta}\log\frac{1}{\delta})n \times \log\log n)$ [17], any set of m disks has complexity $O(m)$ [14], and disks are $(1/4)$-fat, we get the following corollary.

Corollary 3.3 *Let F be a set of n δ-fat triangles and let D be a set of m disks in the plane. Then the complexity of $\mathcal{U}(F \cup D)$ is $O((\frac{1}{\delta}\log\frac{1}{\delta})n\log\log n + \frac{1}{\delta}m)$.*

4 From Locally γ-Fat Objects to Quasi-Triangles

As in most papers on the union complexity of fat objects, we wish to replace our locally γ-fat objects by simpler 'canonical' objects. Let $\mathcal{D}$ be a set of $40/\gamma$ equally spaced orientations, where we assume for simplicity that $40/\gamma$ is an integer. Thus the angle between two consecutive orientations in $\mathcal{D}$ is $\gamma^* := \gamma\pi/20$. We call a direction in $\mathcal{D}$ a *standard direction*. A *quasi-triangle* is an object Δ bounded by two straight edges and one smooth Jordan arc without inflection points.

Definition 4.1 A *γ-standard quasi-triangle* is a quasi-triangle Δ such that:

(a) its two edges have standard directions, and their angle inside Δ is between $\pi - 7\gamma^*$ and $\pi - \gamma^*$;
(b) the tangent line at any point of the Jordan arc makes an angle of at least γ^* with the edges of Δ, and the tangent direction along σ does not vary by more than γ^*.

We say that two γ-standard quasi-triangles are *consistently oriented* if their edges have the same standard orientations.

Fig. 5 illustrates this definition. Observe that property (b) implies that any line parallel to one of the two edges of Δ intersects its Jordan arc at most once. Also note that because of property (a) any set of γ-standard quasi-triangles can be partitioned into $O(1/\gamma)$ subsets of consistently oriented γ-standard quasi-triangles.

The following lemma follows easily from the fact that the angles at each vertex of a γ-standard quasi-triangle are all at least γ^*—this follows from (b)—and that the tangent direction along the arc cannot vary too much.

Fig. 5 The standard directions and four consistently oriented γ-standard quasi-triangles

Lemma 4.2 *A γ-standard quasi-triangle is locally γ'-fat for $\gamma' = \Omega(\gamma)$*

We now set out to reduce the problem of bounding the union of locally γ-fat objects to the problem of bounding the union of γ-standard quasi-triangles. We do this by covering the boundary of each locally γ-fat object using γ-standard quasi-triangles, as follows.

Let o be a locally γ-fat object. We partition ∂o into a number of *subarcs* by putting breakpoints on ∂o in two steps. In the first step we put breakpoints at the following three types of points:

(i) every non-smooth point of ∂o;
(ii) every inflection point of ∂o;
(iii) every smooth point where the tangent line has a standard direction.

Let B_1 be the resulting set of breakpoints. Because o has constant complexity and there are $O(1/\gamma)$ standard directions, $|B_1| = O(1/\gamma)$. In the second step we further refine ∂o by putting breakpoints as follows.

(iv) Put a breakpoint at each point $p \in \partial o$ for which there is a breakpoint $q \in B_1$ such that the line segment pq has a standard direction and $pq \subset o$.

Let B_2 denote the resulting set of breakpoints. We have $|B_2| = |B_1| \cdot O(1/\gamma) = O(1/\gamma^2)$.

Next, we define a γ-standard quasi-triangle $\Delta(\sigma)$ for each of the $O(1/\gamma^2)$ subarcs σ induced by the set of breakpoints. Let p and q be the endpoints of a subarc σ. Assume without loss of generality that pq is parallel to the x-axis and that σ bounds o from above. Let ℓ_p and ℓ_q be the vertical lines through p and q, respectively. Rotate ℓ_p and ℓ_q in counterclockwise direction around p resp. q until they have a standard direction. The γ-standard quasi-triangle $\Delta(\sigma)$ is now formed by σ and two straight segments pr and qr, where r is a point in between ℓ_p and ℓ_q defined as follows. Draw a line ℓ_1 through p whose angle with σ at the point p is $2\gamma^*$. Rotate ℓ_1 clockwise until it reaches a standard direction. Similarly, draw a line ℓ_2 through q whose angle with σ at the point q is $2\gamma^*$, and rotate ℓ_2 counterclockwise until it reaches a standard direction. We define r to be the intersection point of ℓ_1 and ℓ_2—see Fig. 6. Note that the point r must lie below σ and between ℓ_p and ℓ_q.

Lemma 4.3 *The quasi-triangle $\Delta(\sigma)$ formed by σ and the segments pr and qr is a γ-standard quasi-triangle.*

Proof By construction, pr and qr have standard directions. Moreover, the angles that pr and qr make with σ are at least $2\gamma^*$ and at most $3\gamma^*$. Since the tangent direction along σ does not vary by more than γ^*, this implies that the angle between pr and qr is between $\pi - 7\gamma^*$ and $\pi - \gamma^*$, which establishes property (a).

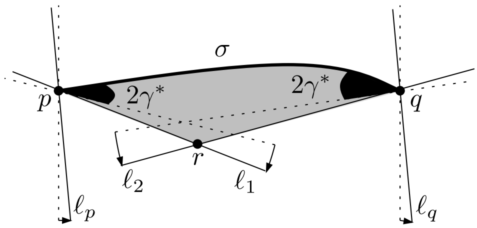

Fig. 6 The γ-standard quasi-triangle $\Delta(\sigma)$ defined for σ

Property (b) follows because the angles that pr and qr make with σ at p resp. q are at least $2\gamma^*$ and the tangent direction along σ does not vary by more than γ^*. □

Lemma 4.4 *The γ-standard quasi-triangle $\Delta(\sigma)$ defined above is contained in o.*

Proof Let x be the lowest point on ℓ_p such that $px \subset o$. Clearly $x \in \partial o$. (Note that it may happen that $x = p$.) Imagine sweeping a segment s from left to right through o, as follows. Start with $s = px$. Move s to the right, keeping it parallel to ℓ_p and keeping its endpoints on ∂o, until s reaches ℓ_q. Note that the upper endpoint of s will move along σ. The lower endpoint of s cannot encounter a breakpoint from B_1 during the sweep, otherwise this breakpoint would have generated a type (iv) breakpoint on σ and σ would not be a subarc. Let σ' be the part of ∂o followed by the lower endpoint of s—see Fig. 7(i). Because there is no breakpoint on σ', we know that σ' does not contain a point where the tangent line has a standard direction. Hence, σ' can cross both pr and qr at most once.

If σ' crosses neither pr nor qr, then $\Delta(\sigma) \subset o$ and we are done, so assume for a contradiction that σ' crosses pr and qr. Take a line through p whose angle with pq is γ^* and a line through q whose angle with pq is γ^*, such that their intersection point r' lies above σ, as in Fig. 7(ii).

Consider the 6-gon defined by the following six points: p, r', q, the intersection of the extension of pr with ℓ_q, r, and the intersection of the extension of qr with ℓ_p—see Fig. 7. Then both σ and σ' are contained in this 6-gon. Let w be the distance between ℓ_p and ℓ_q. Using that angle$(pq, pr') =$ angle$(pq, qr') = \gamma^*$, and that angle$(pq, pr) \leq 4\gamma^*$ and angle$(pq, qr) \leq 4\gamma^*$, one can show that the area of the 6-gon is at most

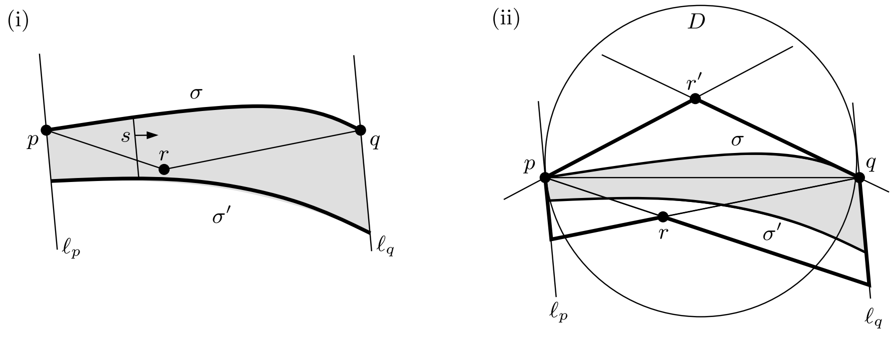

Fig. 7 Illustrations for the proof of Lemma 4.4

$(9/2)\gamma^* w^2$. Now let D be the disk centered at a point of σ and touching ℓ_p and ℓ_q. Because o is locally γ-fat, we have $area(D \sqcap o) \geq \gamma\pi(w\cos(\gamma^*)/2)^2 \geq 0.24\gamma\pi w^2$. On the other hand, $D \sqcap o$ is contained in the area enclosed by σ and σ'. But this area is at most $(9/2)\gamma^* w^2$, which is a contradiction since $\gamma^* = \gamma\pi/20$. □

Now we can reduce the problem of bounding the union complexity of a set of locally γ-fat objects to the problem of bounding the union complexity of a set of locally fat γ-standard quasi-triangles.

Proposition 4.5 *Let $u_\gamma(n)$ denote the maximum complexity of the union of a collection of n consistently oriented γ-standard quasi-triangles. Then the maximum union complexity of any set $\mathcal{F}$ of n locally γ-fat objects is $O((1/\gamma^2) \cdot u_\gamma(n/\gamma^2))$.*

Proof Replace each $o \in \mathcal{F}$ by a collection $T(o)$ of quasi-triangles as described above. Since we have $\partial o \subset \partial\mathcal{U}(T(o))$ and $\mathcal{U}(T(o)) \subset o$ for each object o, the complexity of $\mathcal{U}(\mathcal{F})$ is no more than the complexity of $\mathcal{U}(\{T(o)\colon\ o \in \mathcal{F}\})$.

This gives us a set T of $O(n/\gamma^2)$ γ-standard quasi-triangles. We partition T into $O(1/\gamma)$ subsets T_i of consistently oriented quasi-triangles. Set $n_i := |T_i|$. Every corner of $\mathcal{U}(\mathcal{F})$ will show up either as (i) a corner of $\mathcal{U}(T_i)$ for some i, or (ii) as a corner of $\mathcal{U}(T_i \cup T_j)$ for some pair i, j. The number of corners of type (i) is $\sum_i u_\gamma(n_i)$, which is $O(u_\gamma(n/\gamma^2))$ since $u_\gamma(\cdot)$ is at least linear. To bound the corners of type (ii) we use the Merging Lemma and Lemma 4.2, which imply that the complexity of $\mathcal{U}(T_i \cup T_j)$ is $O((u_\gamma(n_i) + u_\gamma(n_j))/\gamma)$. Hence, the total number of type (ii) corners is bounded by

$$\sum_i \sum_j O\big((u_\gamma(n_i) + u_\gamma(n_j))/\gamma\big)$$
$$= \sum_i \big\{O(1/\gamma) \cdot O(u_\gamma(n_i)/\gamma) + O\big(u_\gamma(n/\gamma^2)/\gamma\big)\big\} = O\big((1/\gamma^2) \cdot u_\gamma(n/\gamma^2)\big). \quad \square$$

5 The Union Complexity of γ-Standard Quasi-Triangles

Let T be a set of consistently oriented γ-standard quasi-triangles. Without loss of generality we assume each $\Delta \in T$ has one edge parallel to the x-axis, and one edge that makes an angle α with the positive x-axis. Recall that $\pi - 7\gamma^* \leq \alpha \leq \pi - \gamma^*$. Set $m := |T|$.

Draw a horizontal line through the horizontal edge and the highest point of every $\Delta \in T$. This partitions the plane into at most $2m + 1$ horizontal strips. Let $\mathcal{T}$ be a balanced binary tree whose leaves correspond to these strips in order. We associate each node ν in $\mathcal{T}$ with the horizontal strip that is the union of the strips corresponding to the leaves in the subtree rooted at ν. We denote this strip by *strip*(ν). Finally, we associate with each node ν a subset $T(\nu) \subset T$, as follows: $\Delta \in T(\nu)$ if Δ completely crosses the strip of ν but does not completely cross the strip of the parent of ν. (This is equivalent to constructing a segment tree [7] on the projections of the Δ's onto the y-axis, and defining $T(\nu)$ to be the quasi-triangles whose projections are stored

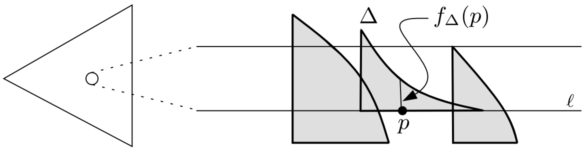

Fig. 8 Illustration for the proof of Lemma 5.1. For clarity the two straight edges of the quasi-triangles are drawn at a right angle (as could be achieved by a suitable transformation) although in fact the angle is almost π

in the canonical subset of ν.) We clip each $\Delta \in T(\nu)$ to *strip*(ν), so that $T(\nu)$ will only contain the parts of these quasi-triangles lying within *strip*(ν). Note that if $N(\Delta)$ is the collection of nodes to which a quasi-triangle Δ is associated, then the clipped pieces of Δ within the strips of the nodes $\nu \in N(\Delta)$ together form Δ. So by assigning Δ to the nodes in $N(\Delta)$ and clipping it, we effectively cut Δ into $|N(\Delta)|$ pieces. Since any Δ is associated to at most two nodes at every level of $\mathcal{T}$—this is a standard property of segment-tree like structures [7]—Δ is cut into $O(\log m)$ pieces. Hence, if we set $m_\nu := |T(\nu)|$ then $\sum_{\nu \in \mathcal{T}} m_\nu = O(m \log m)$.

Lemma 5.1 *The complexity of* $\mathcal{U}(T(\nu))$ *is* $O(\lambda_{s+2}(m_\nu))$.

Proof Let ℓ be the line bounding *strip*(ν) from below. For each $\Delta \in T(\nu)$, we define a function $f_\Delta : \ell \to \mathbb{R}$ as follows. Let $l(p)$ be the line through the point p that makes an angle α with the positive x-axis. Then

$$f_\Delta(p) := \text{ the length of } l(p) \cap \Delta.$$

Recall that the properties of a γ-standard quasi-triangle imply that $l(p)$ intersects Δ in a single segment, if at all. Hence, the boundary of $\mathcal{U}(T(\nu))$ is the upper envelope of the set of functions $\{f_\Delta(p)\colon\ \Delta \in T(\nu)\}$—see Fig. 8—which has complexity $O(\lambda_{s+2}(m_\nu))$ [18], where s is the maximum number of intersections between two object boundaries. □

Next we consider the union of all (clipped) quasi-triangles associated to nodes at a fixed depth in $\mathcal{T}$. Let $N(k)$ denote the nodes of $\mathcal{T}$ at depth k, and define $T(k) := \bigcup_{\nu \in N(k)} T(\nu)$.

Lemma 5.2 *The complexity of* $\mathcal{U}(T(k))$ *is* $O(\lambda_{s+2}(m))$.

Proof The strips of the nodes at a fixed level in the tree are disjoint. Hence, the complexity of $\mathcal{U}(T(k))$ is bounded by the sum of the complexities at each of the nodes $\nu \in N(k)$, which is $\sum_{\nu \in N(k)} O(\lambda_{s+2}(m_\nu))$ by the previous lemma. Since any quasi-triangle can be associated with at most two nodes at any fixed level, we have $\sum_{\nu \in N(k)} m_\nu \le 2m$. The lemma follows. □

To combine the unions of different levels of the tree we use the following lemma.

Lemma 5.3 *The set* $E(k)$ *of boundary edges of* $\mathcal{U}(T(k))$ *has density* $O(1/\gamma^2)$.

Proof For $\nu \in N(k)$, let $E(\nu)$ denote the set of boundary edges of $\mathcal{U}(T(\nu))$. Since the strips of the nodes in $N(k)$ are disjoint, we have $E(k) = \bigcup_{\nu \in N(k)} E(\nu)$.

First we note that, even though the quasi-triangles in $T(\nu)$ are not necessarily fat because they are clipped, the set $E(\nu)$ still has density $O(1/\gamma)$. Indeed, the edges of the union of the unclipped quasi-triangles have density $O(1/\gamma)$ by the Density Lemma; clipping the union to *strip*(ν) can only remove or shorten edges, which does not increase the density.

Now consider an edge $e \in E(\nu)$ and let w be the width of *strip*(ν). We observe that *size*$(e) = O(w/\gamma)$, because at any point $p \in e$ the tangent to e makes an angle $\Omega(\gamma)$ with the boundary lines of the strip. Let D be a disk and assume without loss of generality that D has unit radius. We must argue that D is intersected by $O(1/\gamma^2)$ edges $e \in E(k)$ with *size*$(e) \geq 1$. By the previous observation, we only have to consider edges lying in strips of width $\Omega(\gamma)$. Clearly D can be intersected by only $O(1/\gamma)$ such strips, because the strips in $N(k)$ are disjoint. Since within each strip the density of the boundary edges is $O(1/\gamma)$, the overall density is $O(1/\gamma^2)$. □

We can now prove a bound on the total union complexity of T.

Lemma 5.4 *The complexity of* $\mathcal{U}(T)$ *is* $O((\lambda_{s+2}(m)\log^2 m)/\gamma^2)$.

Proof Consider a corner v of $\mathcal{U}(T)$ that is an intersection point of the boundaries of two quasi-triangles Δ_1 and Δ_2. Let ν_1 be the node such that Δ_1 is associated to ν_1 and the clipped portion of Δ_1 within *strip*(ν_1) contains v. Define ν_2 similarly for Δ_2. Let k_1 and k_2 be the depths of ν_1 and ν_2, respectively. If $k_1 = k_2$ then the corner v is already accounted for in the bound $O(\lambda_{s+2}(m))$ on the complexity of $\mathcal{U}(T(k_1))$. To account for the corners v where $k_1 \neq k_2$ we must consider the unions $\mathcal{U}(T(k_1) \cup T(k_2))$ at different depths k_1 and k_2. Since the sets of boundary edges of $\mathcal{U}(T(k_1))$ and $\mathcal{U}(T(k_2))$ have density $O(1/\gamma^2)$ by the previous lemma, we can use Lemma 5.2 and argue as in the proof of the Merging Lemma: we charge each intersection to the smaller of the two edges, and observe that every edge is charged $O(1/\gamma^2)$ times. This way we can bound the complexity of $\mathcal{U}(T(k_1) \cup T(k_2))$ by $O(\lambda_{s+2}(m)/\gamma^2)$. Because the depth of $\mathcal{T}$ is $O(\log m)$ we now have

$$\sum_{k_1}\sum_{k_2} O\big((\lambda_{s+2}(m))/\gamma^2\big) = O\big((\lambda_{s+2}(m)\log^2 m)/\gamma^2\big),$$

which completes the proof. □

Plugging this result into Proposition 4.5 we get our main theorem.

Theorem 5.5 *The union complexity of any set* $\mathcal{F}$ *of* n *constant-complexity locally* γ*-fat objects is* $O((1/\gamma^6) \cdot \lambda_{s+2}(n)\log^2 n)$.

6 Concluding Remarks

We have introduced a new class of fat objects, namely locally γ-fat objects, and proved that the union complexity of n such objects of constant complexity is

$O((1/\gamma^6) \cdot \lambda_{s+2}(n) \log^2 n)$. This slightly improves the best known bounds on the union complexity of (possibly non-convex and curved) fat objects and extends the result to a more general class of objects. We feel that our definition of fatness is more natural than the existing definitions for non-convex and curved objects: the β-fat objects of van der Stappen et al. [19, 20] can have thin parts (and this may cause the union of n such objects to have $\Omega(n^2)$ complexity), the κ-curved objects of Efrat and Katz [9] are not allowed to have sharp corners (even when the interior angle is large), and the (α, β)-covered objects of Efrat [8] cannot have small parts (e.g. the union of two touching squares is not fat in this definition when the two squares differ a lot in size).

It is unlikely that our bounds are tight: the two logarithmic factors seem artifacts of our proof technique. It would be very interesting to shave off one or two logarithms from the bounds. Note that our results, in particular the Density Lemma, imply that it is sufficient to prove better bounds on the union complexity of so-called consistently oriented fat quasi-triangles.

References

1. Agarwal, P.K., Katz, M.J., Harir, M.: Computing depth orders for fat objects and related problems. Comput. Geom. Theory Appl. **5**, 187–206 (1995)
2. Alt, H., Fleischer, R., Kaufmann, M., Mehlhorn, K., Näher, S., Schirra, S., Uhrig, C.: Approximate motion planning and the complexity of the boundary of the union of simple geometric figures. Algorithmica **8**, 391–406 (1992)
3. Aronov, B., Efrat, A., Koltun, V., Sharir, M.: On the union of κ-curved objects in three and four dimensions. In: Proceedings of the 20th ACM Symposium on Computational Geometry, pp. 383–390 (2004)
4. de Berg, M.: Vertical ray shooting for fat objects. In: Proceedings of the 21st ACM Symposium on Computational Geometry, pp. 288–295 (2005)
5. de Berg, M., Gray, C.: Vertical ray shooting and computing depth orders for fat objects. In: Proceedings of the 17th ACM–SIAM Symposium on Discrete Algorithms (SODA), pp. 494–503 (2006)
6. de Berg, M., Katz, M., van der Stappen, F., Vleugels, J.: Realistic input models for geometric algorithms. Algorithmica **34**, 81–97 (2002)
7. de Berg, M., van Kreveld, M., Overmars, M., Schwarzkopf, O.: Computational Geometry: Algorithms and Applications, 2nd edn. Springer, Berlin (2000)
8. Efrat, A.: The complexity of the union of (α, β)-covered objects. SIAM J. Comput. **34**, 775–787 (2005)
9. Efrat, A., Katz, M.: On the union of κ-curved objects. In: Proceedings of the 14th ACM Symposium on Computational Geometry, pp. 206–213 (1998)
10. Efrat, A., Rote, G., Sharir, M.: On the union of fat wedges and separating a collection of segments by a line. Comput. Geom. Theory Appl. **3**, 277–288 (1993)
11. Efrat, A., Sharir, M.: The complexity of the union of fat objects in the plane. In: Proceedings of the 13th ACM Symposium on Computational Geometry, pp. 104–112 (1997)
12. Katz, M.J.: 3-D vertical ray shooting and 2-D point enclosure, range searching, and arc shooting amidst convex fat objects. Comput. Geom. Theory Appl. **8**, 299–316 (1998)
13. Katz, M.J., Overmars, M., Sharir, M.: Efficient output sensitive hidden surface removal for objects with small union size. Comput. Geom. Theory Appl. **2**, 223–234 (1992)
14. Kedem, K., Livne, R., Pach, J., Sharir, M.: On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles in the plane. Discrete Comput. Geom. **1**, 59–71 (1986)
15. Matoušek, J., Pach, J., Sharir, M., Sifrony, S., Welzl, E.: Fat triangles determine linearly many holes. SIAM J. Comput. **23**, 154–169 (1994)
16. Pach, J., Safruti, I., Sharir, M.: The union of congruent cubes in three dimensions. Discrete Comput. Geom. **30**, 133–160 (2003)
17. Pach, J., Tardos, G.: On the boundary complexity of the union of fat triangles. SIAM J. Comput. **31**, 1745–1760 (2002)

18. Sharir, M., Agarwal, P.K.: Davenport–Schinzel Sequences and Their Geometric Applications. Cambridge University Press, Cambridge (1995)
19. van der Stappen, A.F.: Motion planning amidst fat obstacles. Ph.D. thesis, Utrecht University, Utrecht, the Netherlands (1994)
20. van der Stappen, A.F., Halperin, D., Overmars, M.H.: The complexity of the free space for a robot moving amidst fat obstacles. Comput. Geom. Theory Appl. **3**, 353–373 (1993)
21. van Kreveld, M.: On fat partitioning, fat covering, and the union size of polygons. Comput. Geom. Theory Appl. **9**, 197–210 (1998)