

Fat Triangles Determine Linearly Many Holes[◇]

Jiří Matoušek*
 János Pach **
 Micha Sharir ***
 Shmuel Sifrony [◇]
 Emo Welzl [†]

B 92–13
 June 1992

Abstract

We show that for every fixed $\delta > 0$ the following holds: If F is a union of n triangles, all of whose angles are at least δ , then the complement of F has $O(n)$ connected components, and the boundary of F consists of $O(n \log \log n)$ straight segments (where the constants of proportionality depend on δ). This latter complexity becomes linear if all triangles are of roughly the same size or if they are all infinite wedges.

[◇]Jiří Matoušek and Emo Welzl have been supported by the Deutsche Forschungsgemeinschaft, Schwerpunktprogramm “Datenstrukturen und effiziente Algorithmen”, grant 1265/1-2. János Pach and Micha Sharir have been supported by NSF Grant CCR-89-01484. Emo Welzl and Micha Sharir have also been supported by a grant from the German-Israeli Binational Science Foundation. János Pach has also been supported by Hungarian Science Foundation Grant OTKA-1814. Micha Sharir has also been supported by ONR Grant N00014-90-J-1284, and by grants from the U.S.-Israeli Binational Science Foundation, and the Fund for Basic Research administered by the Israeli Academy of Sciences. The authors wish to thank the DIMACS Center for additional support under National Science Foundation Grant STC-88-09648.

*Department of Computer Science, Charles University, Praha, Malostranské nám. 25, 118 00 Praha 1, Czechoslovakia and Institut für Informatik, Freie Universität Berlin, Arnimallee 2-6, W-1000 Berlin 33.

**Courant Institute of Mathematical Sciences, New York University and Mathematical Institute of the Hungarian Academy of Sciences.

***School of Mathematical Sciences, Tel Aviv University, and Courant Institute of Mathematical Sciences, New York University.

[◇]Courant Institute of Mathematical Sciences, New York University.

[†]Institut für Informatik, FB Mathematik, Freie Universität Berlin, Arnimallee 2-6, D-1000 Berlin 33

1 Introduction

The problem studied in this paper is to obtain sharp upper bounds on the combinatorial complexity of the union of n geometric figures in the plane. This problem arises in many applications. For example, in motion planning for systems with two degrees of freedom, one constructs the two-dimensional configuration space of the system as the complement of the union of n “forbidden regions”, each representing the space of placements of the system in which a collision occurs between two specific system and obstacle features (see [9], [14] for details). It has also been observed recently that families of figures, with the property that the union of any subfamily has small combinatorial complexity, have several additional useful properties. For example, they admit efficient output-sensitive hidden surface removal algorithms (when these figures lie at various heights and are viewed from a point far below them) [20]. Also one can obtain sharp bounds on the number of “ k -sets” in an arrangement of such figures [19] and an efficient algorithm for “point-stabbing” queries in a collection of such figures (where one has to report all figures containing a query point) [19].

The simplest example of a family with the above property is a collection of half planes, each bounded by a line, or more generally by a pseudo-line. A more interesting example is a family of *pseudodisks*, i.e. figures with the property that the boundaries of each pair of them intersect in at most two points. It was shown in [14] that the boundary of the union of n pseudodisks consists of at most $6n - 12$ connected pieces of the boundaries of the given figures (a special case of this result has also been obtained in [8]). Another case was studied in [4], and involved a family of figures, each bounded between a portion of the x -axis and a curve lying above the axis and delimited by two points on the axis, with the property that any pair of these curves intersect in at most 3 points. It was shown that the combinatorial complexity of the union of n such figures is $O(n\alpha(n))$, where $\alpha(n)$ is the inverse Ackermann’s function.

As all these examples indicate, the property of having a union of small combinatorial complexity somehow seems to require that the boundaries of any pair of the given figures intersect in a small number (1,2 or 3) of points. When the allowed number of intersections becomes 4 or more, there are sets of n triangles whose union has quadratic complexity. However, one observes that to attain quadratic complexity, it seems to be essential that the triangles be very narrow and many must have an angle that tends to 0 as n increases.

The purpose of this paper is to show that if this is not allowed, namely if we are given a collection of triangles that are “fat,” then indeed the combinatorial complexity of their union is small.

Statement of results. We call a triangle T δ -fat, if each angle of T is at least δ . By a *figure* we mean a (closed) region in the plane, bounded by a closed Jordan curve or by an unbounded Jordan arc.

Let \mathcal{F} be a finite family of figures. A *hole* of \mathcal{F} is a connected component of the complement of the union of the figures of \mathcal{F} . The number of holes of \mathcal{F} will be

denoted by $H(\mathcal{F})$.

A point of the boundary of the union of a family \mathcal{F} is called a *corner* of \mathcal{F} if it is a point of intersection between the boundaries of two figures in \mathcal{F} . The *boundary complexity* of \mathcal{F} (denoted by $BC(\mathcal{F})$) will be the number of corners of \mathcal{F} ; note that we do not count vertices (if any) of the figures of \mathcal{F} as corners — their number is usually small and presents no problems in the analysis. An *edge* of \mathcal{F} is a connected portion of the boundary of the union of \mathcal{F} contained in the boundary of a single figure between two adjacent corners.

Our main results are the following theorems:

Theorem 1.1 *For any fixed $\delta > 0$, every family \mathcal{F} of n δ -fat triangles has $O(n)$ holes, with the constant of proportionality depending on δ .*

Using this theorem in combination with the Combination Lemma of Edelsbrunner et al in the next section), we will show in Section 4 the following:

Theorem 1.2 *For any fixed $\delta > 0$, the boundary complexity of every family \mathcal{F} of n δ -fat triangles is $O(n \log \log n)$ (again, the constant of proportionality depends on δ). On the other hand, there exist such families (even with $\delta = 60^\circ$) whose boundary complexity is $\Omega(n\alpha(n))$.*

In the special case when the triangles in our family all have roughly the same size, the boundary complexity becomes linear (in the statement of the theorem, $\text{diam}(T)$ denotes the diameter of triangle T):

Theorem 1.3 *Let $\delta > 0$ and $0 < c \leq C$ be fixed numbers. Let \mathcal{F} be a family of n δ -fat triangles, such that $c \leq \text{diam}(T) \leq C$ for every triangle $T \in \mathcal{F}$. Then the boundary complexity of \mathcal{F} is $O(n)$ (with the constant of proportionality depending on δ and on C/c). The boundary complexity is also linear for a family of δ -fat wedges (regions bounded between a pair of rays with a common endpoint).*

Related results have been recently obtained by Alt et al. [2], where the complexity of fat objects was first considered. They showed, among other results, that the boundary complexity of the union of n δ -fat double wedges is $O(n)$. They have also shown that the number of holes (and the boundary complexity) of the union of n triangles, each of which is homothetic either to a fixed triangle T or to the reflection of T , is linear. These results are special cases of the results that we obtain in this paper.

2 Preliminaries

In this section we review two basic results concerning arrangements of certain types of figures, which will be needed in the subsequent analysis. The first result, adapted from [6], is stated here in a more specialized form, which nevertheless follows easily from the original version of [6].

Lemma 2.1 (Combination Lemma [6]) *Let \mathcal{F}_1 and \mathcal{F}_2 be families of figures, whose boundaries are polygons with n_1 and n_2 sides in total. Then*

$$\text{BC}(\mathcal{F}_1 \cup \mathcal{F}_2) \leq \text{BC}(\mathcal{F}_1) + \text{BC}(\mathcal{F}_2) + O(n_1 + n_2 + \text{H}(\mathcal{F}_1 \cup \mathcal{F}_2)).$$

□

The next lemma follows from a more general statement about pseudodisks, [14], [8]. However, for the sake of completeness we present the simple proof for the special case we need here (two figures are called *homothetic*, if one can be obtained from the other by translation and scaling).

Lemma 2.2 *For a family \mathcal{F} of n pairwise homothetic triangles we have $\text{BC}(\mathcal{F}) \leq 6n$.*

Proof: Let us first observe that the boundaries of two homothetic triangles cross in at most two points. Consider now a corner w of \mathcal{F} , which is the intersection of two edges e and e' of two of the triangles. Each edge has one direction at the corner in which the edge ‘disappears’ locally into the respective other triangle. Let v and v' be the vertices incident to the edges in those distinguished directions. Note that either v or v' must be covered by the respective other triangle. Indeed, in order for v to lie outside, the edge e must create another boundary crossing, and similarly for v' ; thus, if both v and v' are not covered, we get at least three boundary crossings, which is impossible. If v is covered, the corner w is the last corner on e in the direction towards v (since, by convexity, the whole portion between w and v is covered); an analogous statement holds for v' .

We charge the corner to the pair (e, v) , if v lies in the other triangle, and to the pair (e', v') , otherwise. We have seen that each such pair can be charged at most once, and so the number of corners is at most twice the number of vertices, namely $6n$. (Note that this bound holds even if we also count in the boundary complexity the triangle vertices on the boundary). □

3 Bounding the Number of Holes

In this section we prove Theorem 1.1, that is, we show that a set of δ -fat triangles has at most a linear number of holes.

Passing to canonical triangles. The first step in the proof is to transform the given collection \mathcal{F} to another collection consisting of *canonical triangles*, so that the number of holes in the new collection is not much different than the number of holes of \mathcal{F} . Specifically, we have:

Lemma 3.1 (Canonization Lemma) *For each $\delta > 0$ there exists a positive constant $c = c(\delta) = O(1/\delta)$, such that if \mathcal{F} is a family of n δ -fat triangles, then there exists*

families $\mathcal{F}_1, \dots, \mathcal{F}_c$ consisting of $O(n)$ triangles in total, such that each \mathcal{F}_i is a family of $\delta/4$ -fat homothetic triangles and

$$H(\mathcal{F}) \leq H(\mathcal{F}_1 \cup \dots \cup \mathcal{F}_c) + O(n).$$

The canonization is achieved by producing triangles which have edges from some fixed finite set of directions $D(\delta) = \{0, \hat{\delta}, 2\hat{\delta}, \dots, (k_\delta - 1)\hat{\delta}\}$, where $k_\delta = \lceil 4\pi/\delta \rceil$ and $\hat{\delta} = 2\pi/k_\delta$. The set $D(\delta)$ has the property that every angle of at least $\delta/2$ contains a direction in $D(\delta)$.

Lemma 3.2 *Let $\delta > 0$. Any δ -fat triangle can be expressed as the union of three $(\delta/2)$ -fat triangles T_1, T_2, T_3 , such that two of the sides of each T_i have directions in $D(\delta)$, while the third is a side of T .*

Proof: Let T be a δ -fat triangle with vertices A, B, C and let O be the center of its inscribed circle (which is also the intersection of the angle bisectors, see Figure 1). Hence each of the angles OAC, OBC is at least $\delta/2$. We can thus find a point Q in the triangle, such that the point O lies in the triangle ABQ , and the segments AQ and BQ have directions in $D(\delta)$. Such a point Q determines the triangle $T_1 = ABQ$, and T_2, T_3 can be constructed in an analogous manner for the two other sides of T . \square




Figure 1: First stage of canonization

In the first stage of canonization, we replace each triangle in \mathcal{F} by three “semi-canonical” triangles as in the preceding lemma. In a second stage we shrink each of the new triangles until it becomes the union of two “fully canonical” triangles. This is shown in the following lemma.

Lemma 3.3 (Shrinking Lemma) *Let \mathcal{F} be a family of n triangles. Let \mathcal{F}^* arise from \mathcal{F} by replacing each triangle $T = ABC$ in \mathcal{F} by the union of two triangles ABX, AYC , such that X lies on AC and Y lies on AB (see Figure 2). Then*

$$H(\mathcal{F}) \leq H(\mathcal{F}^*) + 3n.$$




Figure 2: Second stage of canonization

Proof: Since the union of \mathcal{F}^* is contained in the union of \mathcal{F} , the only way in which the number of holes of \mathcal{F} might decrease as we pass from \mathcal{F} to \mathcal{F}^* is when a pair of holes are merged together to form a single hole. Let us imagine that every triangle ABC of \mathcal{F} shrinks into the corresponding figure of \mathcal{F}^* by a continuous deformation, during which the side BC is deformed into an outward-concave curve γ , e.g. in the manner depicted in Figure 3.




Figure 3: Shrinking a triangle

During this shrinking process, two holes of \mathcal{F} may be merged to form a new hole only when a vertex of some other triangle is passed by γ and appears on the boundary of the union of the shrinking family of figures. Each such event decreases the number of holes by 1, and we can charge this event to the newly appearing vertex. Note that this event is irreversible—once a vertex has appeared on the

boundary of our family, it will never be covered again, so there are at most $3n$ such events during the entire shrinking process, so the number of holes could not have decreased by more than $3n$. \square

Now the proof of Lemma 3.1 is easy. First we replace, using Lemma 3.2, each triangle of the original family \mathcal{F} by the union of a triple of semi-canonical triangles, each having two sides in the set of canonical directions. Then we replace each semi-canonical triangle ABC by a pair of triangles ABX , AYC as in Lemma 3.3, so that each side of the new triangles has a direction in a fixed finite set of directions, and one angle in each triangle is exactly $\hat{\delta}$ (the angle at vertex B and C , respectively); thus the final triangles fall already into a constant number of families of homothetic triangles). We can apply the shrinking of Lemma 3.3 once more to ensure that we have a set of at most $12n$ triangles, where two angles are $\hat{\delta}$. That is, the triangles fall now in $2k_\delta = O(1/\delta)$ homothetic classes. Lemma 3.3 is easily seen to imply that at most $O(n)$ holes can be lost in both shrinking processes, since the number of triangles (and so the number of vertices) is linear. \square

Boundary complexity for a pair of homothetic families By the Canonization Lemma 3.1, it suffices to bound the number of holes of a union of a constant number of families, each consisting of homothetic $\hat{\delta}$ -fat triangles. For simplicity of exposition, we will continue to denote $\hat{\delta}$ by δ . If $\mathcal{F} = \mathcal{F}_1 \cup \dots \cup \mathcal{F}_c$, then any corner of \mathcal{F} must be a corner of some family of the form $\mathcal{F}_i \cup \mathcal{F}_j$, for $1 \leq i, j \leq c$ (this also includes corners that arise within a single family \mathcal{F}_i), thus

$$\text{BC}(\mathcal{F}) \leq \sum_{i,j} \text{BC}(\mathcal{F}_i \cup \mathcal{F}_j).$$

Therefore Theorem 1.1 will be proved if we prove the following:

Lemma 3.4 *Let $\delta > 0$ be fixed. Let \mathcal{F}_1 and \mathcal{F}_2 be families of triangles, each consisting of n δ -fat homothetic triangles. Then $\text{BC}(\mathcal{F}_1 \cup \mathcal{F}_2) = O(n)$ (with a constant of proportionality that depends on δ).*

Proof: Let us put $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$. We will bound the number of edges of the union F of \mathcal{F} . Let us call the edges of the union F_i of \mathcal{F}_i the *superedges* of \mathcal{F}_i , $i = 1, 2$. If e is an edge of F lying on a superedge s of a triangle T of \mathcal{F}_i , we call s the *supporting superedge*, T the *supporting triangle*, and \mathcal{F}_i the *supporting family* of e .

By Lemma 2.2, we know that the boundary complexity of \mathcal{F}_1 and of \mathcal{F}_2 is linear, i.e. the number of superedges is linear in n .

Call an edge e of F *trivial* if e is the first or the last edge of F along its supporting superedge. The number of trivial edges is therefore $O(n)$.

Since edges of F have only six possible directions, it suffices to bound the number of nontrivial edges with one fixed direction. Fix such a direction d , and let e be a nontrivial edge of F having direction d . Suppose e is supported by the family \mathcal{F}_2 . The edge e is adjacent to two edges f and f' , whose respective supporting triangles

T and T' belong to the other family \mathcal{F}_1 , and are thus homothetic. Let s and s' be the supporting superedges of f and f' , respectively.

Call the pair (s, s') an *active pair* of superedges, if they are connected by an edge e as above; we will refer to e as an edge *belonging* to (s, s') .

We claim that the number of active pairs is $O(n)$. Indeed, the superedges of \mathcal{F}_1 are non-intersecting and each active pair is visible from each other in direction d . The number of such visible pairs is linear; this can be seen by sweeping a line in direction d across the plane, or by applying a graph planarity argument.

The proof will therefore be finished if we prove the following:

Lemma 3.5 *Let (s, s') be an active pair of superedges; then the number of non-trivial edges belonging to (s, s') is bounded by a constant (depending on δ).*




Figure 4: Active pair

Proof: Let T and T' be the triangles supporting s and s' , respectively (see Figure 4). Consider all the edges belonging to the active pair (s, s') , which, by our convention, are all assumed to have direction d ; without loss of generality we assume that d is horizontal and that T lies to the left of T' (see Figure 4). Without loss of generality, we may also assume that the corresponding holes of \mathcal{F} lie below these edges.

Let these edges be e_1, \dots, e_m (in ascending order along s and s'). Let b denote the side of T parallel to s' , and let a' denote the side of T' parallel to s . Without loss of generality we may assume that the projection of s in direction d on the line containing b is contained in b (and similarly for a' and s').

For each edge e_i let us denote the intersections of its superedge with the edges b , s' by E_i , E'_i , respectively (see Figure 4; these intersections exist by the assumption just made and because each e_i is non-trivial, and thus penetrates through both T and T'). Consider the parallelogram $E_1 E_m E'_m E'_1$. The angle $E_1 E_m E'_1$ is at least δ , and hence $|E_1 E'_1| \geq \gamma \cdot |E_1 E_m|$, where $\gamma > 0$ is a constant depending on δ .

Consider an edge e_i and its supporting triangle T_i . This triangle must contain both points E_i and E'_i . The key observation is that T_i can not intersect the segment $E_{i+1} E'_{i+1}$, simply because $E_{i+1} E'_{i+1}$ is part of a superedge, and T_i belongs to the same family as $E_{i+1} E'_{i+1}$.

Since T_i is δ -fat, it must contain the triangle R with base $E_i E'_i$ and angles δ at the vertices E_i and E'_i , so R also does not intersect $E_{i+1} E'_{i+1}$. This means that the length of the segment $E_i E_{i+1}$ is at least a constant fraction (depending on δ) of the length of $E_i E'_i$, hence also of $E_1 E_m$. This implies that the number m of nontrivial edges belonging to the active pair (s, s') is bounded by a constant. \square

Remark: A more detailed analysis in the previous lemma shows that the constant claimed is $O(1/\delta^2)$. That is, if we denote the cardinality of \mathcal{F}_i by n_i , then $\text{BC}(\mathcal{F}_1 \cup \mathcal{F}_2) = O(n_i + n_j + \frac{1}{\delta^2} \min\{n_i, n_j\})$. This gives a bound of $O(cn/\delta^2) = O(n/\delta^3)$ for $\text{H}(\mathcal{F})$ (for the original family \mathcal{F}). Summing up, we have at most $O(n/\delta^3)$ holes in the union of n δ -fat triangles. This is probably not tight in terms of δ ; the best lower bound we can derive is $\Omega(n/\delta)$.

In closing this section, we note that Lemma 3.4 has the following corollary, which may be of independent interest. Call a family of triangles *c-oriented* if the orientations of the edges of the triangles are drawn from a fixed set of c orientations; see [10, 11, 18, 22] for several studies of c -oriented polygons.

Corollary 3.6 *The boundary complexity of a family of n c -oriented triangles is $O(n)$, where the constant of proportionality depends on c and the minimum angle between any two of the c given orientations.*

Remark: The weaker result of Alt et al. [2] is also a special case of this corollary.

4 The Boundary Complexity of the Union of Fat Triangles

In this section we analyze the boundary complexity of the union of n fat triangles. We rely on the results of the preceding section concerning the number of holes, on the Combination Lemma 2.1, and on a special way of decomposing the given collection of triangles into subcollections, each having a union with small boundary complexity.

Proof of Theorem 1.2: Let $\text{BC}(n)$ denote the maximum possible boundary complexity of a family of n δ -fat triangles. Let \mathcal{F} be such a family. Applying the first canonization step in the proof of Theorem 1.1, we replace \mathcal{F} by a constant number of subfamilies, each consisting of triangles that have two sides with fixed orientations. By further refining this partitioning, we can also assume that within each subfamily the orientations of the third edges of the triangles all lie within some small angular interval, of length, say 1° . Note that the number of subfamilies is still a constant, and that the overall union of all subfamilies is equal to the union of \mathcal{F} . We will show that the boundary complexity of the union of the triangles in the i -th subfamily is $O(n_i \log \log n_i)$, where n_i is the number of triangles in the subfamily. The Combination Lemma 2.1 then implies that the boundary complexity of the union of all triangles is $O(n \log \log n)$, as asserted.

Thus, from now on, we consider a single subfamily, which, for simplicity, we also denote by \mathcal{F} . By applying an appropriate affine transformation, we can assume that each triangle is a right triangle with one horizontal edge and one vertical edge, that these edges meet in the lower-left vertex of the triangle, and that the hypotenuse of the triangle has orientation between, say 134 and 136 degrees (so the triangle is nearly isosceles).

Our first step is to partition \mathcal{F} into $O(\log n)$ subfamilies so that the boundary complexity of each subfamily is almost linear in the number of triangles it contains.

Lemma 4.1 *If all triangles in \mathcal{F} have the form assumed above, and meet a common horizontal line, then $\text{BC}(\mathcal{F}) = O(n \cdot 2^{\alpha(n)})$, where $\alpha(n)$ is the inverse Ackermann's function.*

Proof: Without loss of generality, assume the line is the x -axis. For each triangle $T \in \mathcal{F}$ let T^+ denote its portion above the x -axis, and T^- denote its portion below the x -axis. The boundary complexity of \mathcal{F} is clearly bounded by the sum of the boundary complexities of the union of the triangles T^+ and of the union of the trapezoids T^- . The boundary complexity of the upper triangles T^+ is $O(n)$: if we direct all edges towards the x -axis, then, as is easily seen, every corner is the last corner (in this direction) for one of its two edges.

As to the lower trapezoids T^- , we first decompose each T^- into two interior-disjoint portions, one being an axis-parallel rectangle and the other being a right, nearly isosceles triangle with a horizontal edge and a vertical edge whose top vertex lies on the x -axis; see Figure 5.

It suffices to show that the boundary complexity of the union of the family \mathcal{F}^- consisting of these new triangles is $O(n \cdot 2^{\alpha(n)})$, because the boundary complexity of the union of the rectangles is trivially linear and the Combination Lemma 2.1 implies that merging the rectangles with the triangles of \mathcal{F}^- yields a joint boundary complexity that is proportional to the complexity of \mathcal{F}^- . We therefore restrict our attention only to the union of \mathcal{F}^- .

Claim 1: A horizontal edge of a triangle T in \mathcal{F}^- can be incident to at most four hole corners.




Figure 5: A triangle cut by a horizontal line

Proof: Let e be the given edge, and let $e' = XY$ be an interval along e that appears on the boundary of the union, and is not the leftmost such interval along e . The left endpoint X of e' is the intersection of e with the hypotenuse of another triangle T' , and our assumption concerning e' implies that the vertical edge of T' also cuts e . See Figure 6. Let the top angles of T, T' be α, α' respectively, and let the length of the vertical edge of T be h ; let g denote the intersection $e \cap T'$.




Figure 6: Two ‘interleaving’ triangles

We have $|e| = h \tan \alpha$ and $|g| = h \tan \alpha'$. Thus $|g|/|e| = \tan \alpha' / \tan \alpha$ is very close to 1, in particular it is greater than $1/2$. This shows that the interval e' is unique, so that e can contain at most two intervals that bound holes, namely e' and another leftmost interval. This completes the proof of the claim. \square

Claim 2: The total number of hole corners that are incident to either a horizontal edge or to a vertical edge is $O(n)$.

Proof: Claim 1 implies that the number of hole corners along horizontal edges is $O(n)$. Let e be a vertical edge and let e' be an interval along e bounding a hole. It is easily verified that the top endpoint of e' must be incident to a horizontal edge. The claim is now immediate. \square

It therefore remains to consider only hole corners formed by intersections of two

hypotenuses of the triangles of \mathcal{F}^- . We order these corners in lexicographical order, so that $c_1 \prec c_2$ if, for $c_1 = (x_1, y_1)$ and $c_2 = (x_2, y_2)$, either $x_1 < x_2$ or $x_1 = x_2$ and $y_1 < y_2$. This is clearly a linear order.

Our strategy is to transform this sequence of corners to a *Davenport Schinzel sequence* of order 4 [1, 12], which will then yield the asserted bound on the boundary complexity of \mathcal{F}^- . (Recall that a Davenport Schinzel sequence of order 4 is a sequence that does not have any two equal adjacent elements, and does not contain as a (not necessarily contiguous) subsequence an alternation $a \cdots b \cdots a \cdots b \cdots a \cdots b$ of length 6 between any two distinct symbols a and b .) To this end, we divide each hypotenuse at its midpoint into two subsegments of equal length, which we refer to as its top part and bottom part, respectively. For each corner c consider the hypotenuse incident to c and appearing along the hole just below c ; c is *associated* with the part (top or bottom) of that hypotenuse, to which it is incident. See Figure 7.




Figure 7: The corner c is associated with the bottom part of the triangle T

We proceed through the ordered sequence of corners and form a sequence U , consisting of all associated appearances of the top or bottom hypotenuse parts in the order that the corresponding corners are encountered. Thus U is composed of at most $2n$ distinct symbols.

Claim 3: The number of appearances of bottom parts in U is at most n , and the number of pairs of equal consecutive elements in U is $O(n)$.

Proof: We first show that no bottom part of a hypotenuse can appear twice in U . Indeed, let T be a triangle with a hypotenuse h and let c be a hole corner of the kind we consider that is associated with the bottom part of h . Thus there exists another triangle T' whose hypotenuse meets T at c and has a smaller slope than h . A calculation similar to that in the proof of Claim 1 shows that the next higher appearance of h along a hole must already appear on its top part. This establishes the first assertion of the claim.

Next consider adjacent equal elements of U . Suppose a hypotenuse h of some triangle T appears twice consecutively in U . Thus h contains two subintervals e , e' that bound holes. But then the bottom endpoint c of the higher of these two intervals must be incident to a vertical edge (otherwise c is incident to some other

hypotenuse h' , which necessarily appears in U between the two appearances of h). The claim is thus an immediate consequence of Claim 2. \square

We can therefore delete from U all bottom appearances and then delete one of each pair of equal consecutive elements. The new sequence U^* consists of only top hypotenuse parts (so it is composed of at most n distinct symbols), has no pair of equal adjacent elements, and satisfies $|U| = |U^*| + O(n)$.

We claim that U^* is indeed a Davenport Schinzel sequence of order 4. That is, we have to show that U^* cannot contain an alternating subsequence of the form $a \cdots b \cdots a \cdots b \cdots a \cdots b$, where a and b are top parts of the hypotenuses of two distinct respective triangles, T, R .

Suppose to the contrary that such an alternation exists. We distinguish between two cases:

I. T and R intersect in at most two points. This can happen in one of the four schematic forms shown in Figure 8.




Figure 8: Two triangles of \mathcal{F}^- intersecting in at most 2 points

Cases (i) and (iv) are easy, because they allow no alternation of a and b in U , as is easily checked. In case (ii) let us first assume that a is the top part of the hypotenuse of the left triangle. Note that all appearances of a between the first and last appearances of b correspond to corners that lie in the vertical strip spanned by the right triangle R . Let p and q be two subintervals of a that give rise to two such appearances of a . Then it is easily seen that there must exist another triangle Q that cuts the hypotenuse of T in some interval between p and q ; see Figure 9. Denote the top angles of triangles T, R, Q by $\alpha, \alpha', \alpha''$, respectively. Let d' denote the length of the vertical edge of R , let d_0 denote the vertical distance between the bottom endpoint of p and the top endpoint of q , and let d denote the vertical distance from the top endpoint of q to the x -axis; see Figure 9.

Simple trigonometric calculations show that

$$d \tan \alpha'' = d_0 \tan \alpha < d' \tan \alpha'$$

and

$$d_0 + d' < d,$$

which is clearly impossible, since all three angles $\alpha, \alpha', \alpha''$ are close to 45° .

This argument implies that, between the first and last appearance in U^* of the hypotenuse of the right triangle, there can be at most one appearance of the hypotenuse of the left triangle. Thus the maximum length of an alternation between a and b in




Figure 9: Case I(ii) of the proof

U^* is 5 (in the form $a \cdots b \cdots a \cdots b \cdots a$). If a is the top part of the hypotenus of the right triangle, the above analysis shows that the longest possible alternation is now only $a \cdots b \cdots a \cdots b$. Exactly the same analysis applies in case (iii).

II. T and R intersect in four points. This is depicted in Figure 10. Again without loss of generality we can assume that T is the triangle whose top vertex lies to the left of that of R (otherwise, as above, the maximum possible alternation will be shorter).




Figure 10: Case II of the proof

Note that the second appearance of a in the alternation must be to the right of the intersection point of the two hypotenuses, which implies that the two last appearances of b in the alternation must occur below the horizontal edge of T . But then, arguing as in the proof of Claim 1, it is easy to show that the last occurrence of a in the alternation must be at the bottom part of the hypotenus, contrary to assumption. Thus the alternation is impossible.

Hence U^* is indeed a Davenport Schinzel sequence of order 4 composed of at most n distinct symbols, so its length is at most $O(n \cdot 2^{\alpha(n)})$ [1]. This is also an upper bound on the length of U , and this clearly completes the proof of the lemma.

□

We now decompose \mathcal{F} as follows. We first find a horizontal line ℓ with the property that the number of triangles in \mathcal{F} lying completely above ℓ and the number of triangles lying completely below ℓ are both at most $n/2$. Let \mathcal{F}_1 denote the subfamily of triangles of \mathcal{F} intersecting ℓ . We apply the same procedure to the two subfamilies of \mathcal{F} consisting respectively of the triangles lying above ℓ and of those lying below ℓ . For each of these subfamilies we find a “halving” horizontal line as above, and define \mathcal{F}_2 to be the collection of triangles in these subfamilies which intersect one of these halving lines. We are now left with four subfamilies, each of which is next halved by a line, and \mathcal{F}_3 consists of the remaining triangles that

intersect one of these lines. Continuing in this fashion, we obtain a decomposition of \mathcal{F} into $O(\log n)$ subfamilies, $\mathcal{F}_1, \mathcal{F}_2, \dots$, and the preceding lemma is easily seen to imply that the boundary complexity of each subfamily \mathcal{F}_i is $O(n_i \cdot 2^{\alpha(n_i)})$, where $n_i = |\mathcal{F}_i|$.

We now apply the Combination Lemma 2.1 in a tree-like fashion. That is, we merge the subfamilies \mathcal{F}_i two at a time, then merge each of the resulting collections two at a time, and so on, until all subfamilies are merged together. At each step, when merging two subfamilies $\mathcal{G}_1, \mathcal{G}_2$ to form a combined subfamily \mathcal{G} , we have

$$\text{BC}(\mathcal{G}) \leq \text{BC}(\mathcal{G}_1) + \text{BC}(\mathcal{G}_2) + O(n_1 + n_2) ,$$

where n_i is the size of \mathcal{G}_i , for $i = 1, 2$. This is an immediate consequence of the Combination Lemma and of the fact that the number of holes of \mathcal{G} is $O(n_1 + n_2)$. Since the depth of the tree representing these merges is $O(\log \log n)$ and the sum of the boundary complexities of the individual subfamilies \mathcal{F}_i is $O(n \cdot 2^{\alpha(n)})$, it follows easily that $\text{BC}(\mathcal{F}) = O(n \log \log n)$.

To obtain the lower bound in Theorem 1.2, take a collection of n line segments whose lower envelope consists of $\Omega(n\alpha(n))$ subsegments [23]. Flatten the collection in the y -direction until all segments have almost horizontal slope. Then replace each segment e by an equilateral triangle lying above e and having e as one of its sides. It is easily checked that the boundary complexity of the union of these triangles is $\Omega(n\alpha(n))$. \square

Remark: By modifying the above lower bound construction, and exploiting the special structure of the construction in [23], one can also obtain a collection of n equilateral triangles, whose union has $\Theta(n)$ holes, so that no triangle appears more than once along the boundary of any single hole, and yet the overall boundary complexity is $\Omega(n\alpha(n))$.

Proof of Theorem 1.3: Recall that the theorem asserts that if \mathcal{F} is a family of n δ -fat triangles with $c \leq \text{diam}(T) \leq C$ for all $T \in \mathcal{F}$, then $\text{BC}(\mathcal{F}) = O(n)$, with the constant of proportionality depending on δ and C/c . Let \mathcal{F} be such a family of triangles. We choose a real number D that satisfies the following two conditions:

- (i) No square with side D is intersected by more than two sides of any triangle from \mathcal{F} .
- (ii) The diameter of any triangle of \mathcal{F} is at most a constant multiple of D .

The existence of such a D is guaranteed by the assumptions on \mathcal{F} ; the constant factor in condition (ii) is easily seen to be of the form $\frac{C}{c}\phi(\delta)$, for an appropriate function ϕ .

Let us cover the plane by a grid of squares with side length D . By the choice of D , every triangle of \mathcal{F} intersects at most a constant number of squares of this grid. We claim that the boundary complexity of \mathcal{F} inside each grid square is linear in the number of triangles intersecting that square, and this will imply that the total boundary complexity of \mathcal{F} is $O(n)$.

Let us consider a fixed square Q of the grid. For each triangle $T \in \mathcal{F}$, at most two sides of T intersect Q , hence there exists a wedge (angle) W_T such that $Q \cap W_T = Q \cap T$. Let us consider the family

$$\mathcal{W} = \{W_T : T \in \mathcal{F}, T \cap Q \neq \emptyset\}.$$

The boundary complexity of \mathcal{W} is an upper bound for the complexity of the part of the boundary of \mathcal{F} inside Q . Adapting Theorem 1.1 to the special case of wedges, it is easily seen that the family \mathcal{W} has a linear number of holes. We claim that \mathcal{W} can be partitioned into a constant number of subfamilies $\mathcal{W}_1, \dots, \mathcal{W}_c$, each of which has a linear boundary complexity. Applying the Combination Lemma 2.1 (as in the preceding proof) a constant number of times, we obtain a linear bound on the boundary complexity of \mathcal{W} . (This part of the proof also establishes the second assertion in Theorem 1.3 concerning the complexity of the union of fat wedges.)

We may assume that the apex angle of each wedge of \mathcal{W} is at least δ (this is obvious for triangles having two sides intersecting Q ; for triangles with only one intersecting side, the choice of the apex and its angle are fairly arbitrary). It follows that there exists a fixed set of a constant number $c = O(1/\delta)$ of canonical orientations (e.g. $\delta/2$ apart from each other) so that each wedge in \mathcal{W} contains a ray emerging from its apex and having one of these canonical orientations. We thus choose the decomposition $\mathcal{W} = \mathcal{W}_1 \cup \dots \cup \mathcal{W}_c$ so that for all wedges in the same subfamily, the corresponding rays are all in the same (canonical) direction. It is well-known that the boundary complexity of each subfamily \mathcal{W}_i of wedges is linear. Indeed, if the common ray direction is assumed to be the negative y -direction, the boundary of the union of \mathcal{W}_i is the upper envelope of the collection of rays that bound these wedges, and it is known that the complexity of such an envelope is linear (see e.g. [5]). This finishes the proof of Theorem 1.3. \square

5 Extensions, Applications, and Open Problems

We have so far shown that the union of n δ -fat triangles has a linear number of holes and that its boundary complexity is $O(n \log \log n)$, and can be $\Omega(n\alpha(n))$. In this section we consider several extensions of these results, mention some applications, and conclude with some open problems.

Constructing the union of fat triangles. First we note that one can also calculate efficiently the union of such a family \mathcal{F} . The following algorithm, adapted from [14], can be used. Partition \mathcal{F} into two subfamilies of roughly $n/2$ triangles each. Recursively calculate the union F_1 of \mathcal{F}_1 and the union F_2 of \mathcal{F}_2 . Then merge the two unions by the line sweeping procedure of Chazelle and Edelsbrunner [3] or of Mairson and Stolfi [15]. This computes all k intersections between the boundaries of F_1 and of F_2 in time $O(N \log N + k)$, where N is the overall size of F_1 and of F_2 . But each such intersection is easily seen to be a corner of the overall union of \mathcal{F} , so by Theorem 1.2 we have that both N and k are bounded by

$O(n \log \log n)$. This easily implies that we can construct the union of \mathcal{F} from F_1 and F_2 in time $O(n \log n \log \log n)$, so the overall running time of this algorithm is $O(n \log^2 n \log \log n)$. We thus have

Theorem 5.1 *One can calculate the union of n δ -fat triangles in $O(n \log^2 n \log \log n)$ time and $O(n \log \log n)$ storage (where the constant of proportionality depends on δ).*

Remark: One should contrast the problem of explicit construction of the union of a collection of figures to that of computing various measures of the union, such as its area or the length of its boundary. Such measures can be calculated efficiently for the case of axis-parallel rectangles, not necessarily δ -fat [17]. However, such efficient procedures are not known for general non δ -fat collections. For δ -fat collections they are immediate by-products of the algorithm given above.

Recently, after the original submission of the paper, Miller and Sharir [16] obtained an improved randomized incremental algorithm for computing the union of n fat triangles, using $O(n \cdot 2^{\alpha(n)} \log n)$ expected time and storage.

General “fat” objects We can also extend our results to families of polygons, which can be expressed as the union of a constant number of δ -fat triangles. Some “fatness” condition is clearly essential for such a result to hold, since one can form a quadratic number of holes with very narrow objects. Moreover, the following example shows that even when the polygons appear to be fat in an intuitive sense, they can still form quadratically many holes, so a stronger condition, such as imposed above, has to be enforced.

Example 5.2 *There exists a family of n similar convex figures (actually regular polygons), for each of which the ratio between the radii of the circumscribed and inscribed circles is less than 2, and which determine $\Omega(n^2)$ holes.*

Proof: We will construct a family of $2n$ regular n -gons. Let us choose a regular n -gon $A = A_1 A_2 \dots A_n$. On each of its sides, $A_i A_{i+1}$, choose $n + 1$ equidistant points $B_{i,0} = A_i, B_{i,1}, \dots, B_{i,n-1}, B_{i,n} = A_{i+1}$. The first half of our family consists of n regular n -gons C_1, C_2, \dots, C_n , where $C_j = B_{1,j} B_{2,j} \dots B_{n,j}$. The second half of the family consists of n regular n -gons D_1, \dots, D_n , where D_i arises as a mirror image of A , reflected around the side $A_i A_{i+1}$. This family determines quadratically many holes. \square

Remarks. (1) This example is somewhat misleading, because we ignore here the overall description complexity of the polygons C_i, D_i (which is itself quadratic). We include this example only to demonstrate that one needs to be careful in the definition of fatness if one wishes to extend the results of this paper to more complex figures than triangles.

(2) The reason for the large complexity in this example is that the boundaries of the convex figures intersect in many points per pair. It remains to investigate what happens if we consider a family of fat objects, such that the number of intersections of boundaries of any pair is bounded by a constant.

Applications. As briefly mentioned in the introduction, the fact that the boundary complexity of a family of fat triangles is small has various combinatorial and algorithmic consequences. So far these applications were limited to the case of pseudodisks and to a few other favorable cases mentioned in the introduction. These applications can now be extended to the case of fat triangles. We list some of them as corollaries of the bounds obtained in the preceding sections, and omit the proofs, which are easily obtained by adapting the earlier proofs cited below.

Corollary 5.3 *Let T_1, \dots, T_n be n δ -fat triangles lying in three dimensional space in arbitrary horizontal planes and viewed from a point at $z = -\infty$. Then one can perform hidden surface removal for this scene in time $O(n^{3/2} \log n (\log \log n)^{1/2} + k)$, where k is the size of the resulting “visibility map”.*

Proof: See [20]. □

Remark: Recently, after the original submission of this paper, this result has been significantly improved in [13]. The algorithm presented there is also based on the results of this paper, and its running time is $O((n \log \log n + k) \log^2 n)$.

Corollary 5.4 *Let T_1, \dots, T_n be n δ -fat triangles in the plane, and let $k \leq n - 2$ be an integer. The number of intersection points of the boundaries of these triangles which are covered by at most k other triangles is $O(nk \log \log \frac{n}{k})$.*

Proof: See [19]. □

Corollary 5.5 *Let T_1, \dots, T_n be n δ -fat triangles in the plane. One can preprocess them by a randomized algorithm, whose expected running time is $O(n \log^2 n \log \log n)$, into a data structure of size $O(n \log n \log \log n)$, so that, given any query point z , all k triangles containing z can be reported in (worst-case) time $O((k + 1) \log n)$.*

Proof: See [19]. □

Remark: The bounds stated in the preceding theorems follow from the bound $O(n \log \log n)$ on the boundary complexity of a collection of fat triangles. Since we believe that this bound is not tight (see below), we expect corresponding improvements in the bounds of the preceding theorems. We also note that the running time of the algorithm of Corollary 5.5 can be slightly improved by the recent technique of [16] mentioned above.

Recently, the results of this paper have been applied in [21] to obtain efficient algorithms for motion planning among fat obstacles.

Open problems. The main open problem that arises is to close the gap between the upper and lower bounds on the maximum boundary complexity of a union of n δ -fat triangles. We venture the conjecture that the correct bound is indeed $O(n\alpha(n))$. It is annoying that we were unable to prove this even in the special case of Lemma 4.1.

Another problem is to calibrate the dependence of the constants of proportionality in the various bounds obtained above in terms of δ , so that sharp bounds can be obtained also in cases where δ does depend on n . Some progress towards this goal was recently achieved in [7].

Finally, it is challenging to extend the results of this paper to three dimensions. For example, can one show that the boundary complexity of the union of n arbitrary cubes (or of ‘fat’ simplices) in 3-space is close to *quadratic* in n (as opposed to a trivial cubic upper bound)?

References

- [1] P. AGARWAL, M. SHARIR AND P. SHOR, Sharp upper and lower bounds on the length of general Davenport Schinzel sequences, *J. Combin. Theory, Ser. A* 52 (1989), pp. 228–274.
- [2] H. ALT, R. FLEISCHER, M. KAUFMANN, K. MEHLHORN, S. NÄHER, S. SCHIRRA AND C. UHRIG, Approximate motion planning and the complexity of the boundary of the union of simple geometric figures, *Proc. 6th Symp. on Computational Geometry*, 1990, pp. 281–289.
- [3] B. CHAZELLE AND H. EDELSBRUNNER, An optimal algorithm for intersecting line segments in the plane, *J. ACM* 39 (1992), 1–54.
- [4] H. EDELSBRUNNER, L. GUIBAS, J. HERSHBERGER, J. PACH, R. POLLACK, R. SEIDEL, M. SHARIR AND J. SNOEYINK, On arrangements of Jordan arcs with three intersections per pair, *Discrete Comput. Geom.* 4 (1989), pp. 523–539.
- [5] H. EDELSBRUNNER, L. GUIBAS, J. PACH, R. POLLACK, R. SEIDEL AND M. SHARIR, Arrangements of curves in the plane – Topology, combinatorics, and algorithms, *Theoretical Computer Science* 92 (1992), 319–336.
- [6] H. EDELSBRUNNER, L. GUIBAS AND M. SHARIR, The complexity and construction of many faces in arrangements of lines and of segments, *Discrete Comput. Geom.* 5 (1990), pp. 161–196.
- [7] A. EFRAT, G. ROTE AND M. SHARIR, On the union of fat wedges and separating a collection of segments by a line, *Computational Geometry — Theory and Applications*. manuscript, 1992.
- [8] S. FORTUNE, A fast algorithm for polygon containment by translation, *Proc. 12th Int. Colloq. on Automata, Languages and Programming*, 1985, pp. 189–198.
- [9] L. GUIBAS, M. SHARIR AND S. SIFRONY, On the general motion planning problem with two degrees of freedom, *Discrete Comput. Geom.* 4 (1989), pp. 491–521.

- [10] R.H. GÜTING, *Conquering Contours: Efficient Algorithms for Computational Geometry*, Ph.D. Dissertation, Dortmund University, 1983.
- [11] R.H. GÜTING AND TH. OTTMANN, New algorithms for special cases of hidden line elimination problems, *Computer Vision, Graphics, and Image Processing* 40 (1987), 188–204.
- [12] S. HART AND M. SHARIR, Nonlinearity of Davenport Schinzel sequences and of generalized path compression schemes, *Combinatorica* 6 (1986), pp. 151–177.
- [13] M. KATZ, M. OVERMARS AND M. SHARIR, Efficient output sensitive hidden surface removal for objects with small union size, to appear in *Computational Geometry — Theory and Applications*.
- [14] 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), pp. 59–71.
- [15] H. MAIRSON AND J. STOLFI, Reporting and counting intersections between two sets of line segments, in *Theoretical Foundations of Computer Graphics and CAD* (R.A. Earnshaw, ed.), NATO ASI Series, Vol. F-40, Springer Verlag, Berlin, 1988, pp. 307–325.
- [16] N. MILLER AND M. SHARIR, Efficient randomized algorithms for constructing the union of fat triangles and of pseudodiscs, manuscript, 1991.
- [17] F. PREPARATA AND M. SHAMOS, *Computational Geometry – An Introduction*, Springer Verlag, Berlin 1985.
- [18] G. RAWLINS AND D. WOOD, Optimal computation of finitely oriented convex hulls, *Information and Computation* 72 (1987), 150–166.
- [19] M. SHARIR, On k -sets in arrangements of curves and surfaces, *Discrete Comput. Geom.* 6 (1991), 593–613.
- [20] M. SHARIR AND M. OVERMARS, A simple output sensitive hidden surface removal algorithm, *ACM Trans. Graphics* 11 (1992), 1–11.
- [21] A.F. VAN DER STAPPEN, D. HALPERIN AND M. OVERMARS, The complexity of free space for a robot moving amidst fat obstacles, manuscript, 1992.
- [22] X.H. TAN, T. HIRATA AND Y. INAGAKI, The intersection searching problem for c -oriented polygons, *Inform. Proc. Lett.* 37 (1991), 201–204.
- [23] A. WIERNIK AND M. SHARIR, Planar realization of nonlinear Davenport Schinzel sequences by segments, *Discrete Comput. Geom.* 3 (1988), pp. 15–47.