Improved Bound for the Union of Fat Triangles*

Esther Ezra† Boris Aronov‡ Micha Sharir§

July 7, 2010

Abstract

We show that, for any fixed $\delta > 0$, the combinatorial complexity of the union of n triangles in the plane, each of whose angles is at least δ, is $O(n2^{\alpha(n)} \log^* n)$, with constant of proportionality depending on δ. This considerably improves the twenty-year-old bound $O(n \log\log n)$, due to Matoušek *et al.* [29, 30].

*Work on this paper by Boris Aronov and Micha Sharir has been supported by a joint grant No. 2006/194 from the U.S.-Israel Binational Science Foundation. Work by Boris Aronov has also been supported by NSA MSP Grant H98230-10-1-0210 and NSF Grant CCF-08-30691. Work by Micha Sharir has also been supported by NSF Grant CCF-08-30272, by Grant 338/09 from the Israel Science Fund, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

†Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; *esther@courant.nyu.edu*.

‡Department of Computer Science and Engineering, Polytechnic Institute of NYU, Brooklyn, NY 11201-3840, USA; *aronov@poly.edu*.

§School of Computer Science, Tel Aviv University, Tel Aviv 69978 Israel and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; *michas@post.tau.ac.il*.

1 Introduction

Let $\mathcal{T}$ be a collection of n triangles in the plane. Let $\mathcal{A}(\mathcal{T})$ denote the two-dimensional *arrangement* induced by the edges of the triangles in $\mathcal{T}$. This is the decomposition of the plane into *vertices*, *edges*, and *faces* (also called *cells*), each being a maximal connected set contained in the intersection of some fixed subset of the edges of the triangles of $\mathcal{T}$ and not meeting any other edge. To simplify the presentation, and with no real loss of generality, we assume that the triangles of $\mathcal{T}$ are in *general position*, so that, in paticular, no point of the plane is incident to more than two edges of the triangles.

The *combinatorial complexity* of the union of the triangles in $\mathcal{T}$ is the number of vertices and edges of the arrangement appearing along the union boundary. In this paper we study the union complexity in the special case where all the given triangles are *δ-fat* for some fixed $\delta > 0$. That is, each angle of every triangle is at least δ.

Previous results. The problem of bounding the combinatorial complexity of the union of geometric objects has received considerable attention in the past 25 years. See [2] for a recent comprehensive survey of the area.

The case involving *pseudodiscs* (that is, a collection of simply-connected planar regions, where each pair of their boundaries intersect at most twice) arises for Minkowski sums of a fixed convex object with a set of pairwise disjoint convex objects (which is the problem one faces in translational motion planning of a convex robot in the plane), and has been studied by Kedem *et al.* [25], who showed that in this case the union has only linear complexity.

In general, the complexity of the union of n simply-shaped objects in the plane (of triangles, say) can be $\Theta(n^2)$. However, as shown in a series of papers during the past 20 years, the bound becomes near-linear when the objects are "fat." Efrat *et al.* [16] showed that the union of n unbounded *δ-fat wedges* in the plane (that is, each of their opening angles is at least δ) is $O(n)$, with a constant of proportionality that depends on δ. Later, Matoušek *et al.* [29,30] proved that the union of n *δ-fat triangles* in the plane has only $O(n)$ *holes* (connected components of its complement), and that its combinatorial complexity is $O(n \log\log n)$. The constant of proportionality, which depends on the fatness factor δ, was later improved by Pach and Tardos [31]. The work by Matoušek *et al.* [29,30] has been generalized in various ways. In particular, there have been several recent studies involving fat objects in $\mathbb{R}^3$; see, e.g., [19,21] and the references therein. Other studies of the planar case extend the analysis to the realm of curved fat objects. Efrat and Sharir [17] studied the union of planar fat convex objects. Here we say that a planar convex object c is *γ-fat*, for some fixed $\gamma > 1$, if there exist two concentric disks, $D \subseteq c \subseteq D'$, such that the ratio between the radii of D' and D is at most γ. In this case, the combinatorial complexity of the union of n such objects, such that the boundaries of each pair of objects intersect in a constant number of points, is $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$. With an appropriate (stronger) definition of fatness, some of the previous work extends to non-convex objects, resulting in sharper nearly-linear bounds; see Efrat and Katz [15], Efrat [14] and the very recent work of De Berg [12] (improving the bound previously established in [11]).

Still, the case of fat triangles can be regarded as the "flagship" problem in this collection, being among the simplest and most natural instances, and yet posing challenges, of obtaining sharper bounds on the union complexity, which have not been successfully addressed for nearly twenty years.

Our result. The best known lower bound for the maximum complexity of the union of n δ-fat triangles is $\Omega(n\alpha(n))$, where $\alpha(\cdot)$ is the extremely slowly growing inverse Ackermann function [32]. This bound is obtained by constructing a *lower envelope* of n segments of complexity $\Omega(n\alpha(n))$

(see, e.g., [32] for the construction), sufficiently flattening the construction in the y-direction, and then extending each segment upwards into an equilateral (and thus fat) triangle. The prevailing conjecture, originally posed by Matoušek *et al.* [30], is that the maximum complexity is $\Theta(n\alpha(n))$.

In this paper we improve the $O(n \log\log n)$ upper bound established in [30] to $O(n2^{\alpha(n)} \log^* n)$, thereby making considerable progress towards settling this conjecture.

Besides being a significant advance on a problem which is interesting and challenging in its own right, our result has several useful applications, all obtained by plugging the new bound into the appropriate existing machinery. They include (see Section 3 for details):
(i) A $O(\log\log^* \text{OPT})$ approximation factor for the *set cover* problem for δ-fat triangles and points in the plane, where OPT is the size of the smallest set cover. This improves the factor $O(\log\log\log \text{OPT})$ previously obtained in [3].
(ii) A $O((1/\varepsilon)\log\log^*(1/\varepsilon))$ bound on the size of an ε-net for the "dual" range space involving fat triangles and points in the plane, improving the previous bound $O((1/\varepsilon)\log\log\log(1/\varepsilon))$ derived in [3].
(iii) An improved bound on the running time of the output-sensitive algorithm of [24] for *hidden surface removal* for horizontal δ-fat triangles in 3-space. With our new bound on the union complexity, this algorithm can compute the visibility map of the scene as seen from a given view point, in time $O((n2^{\alpha(n)} \log^* n + k)\log^2 n)$ and $O(n2^{\alpha(n)} \log^* n \log n)$ storage, where k is the output size. This is an improvement over the previous respective bounds $O((n \log\log n + k)\log^2 n)$, $O(n \log\log n \log n)$ obtained in [24].

2 Bounding the Union Complexity

2.1 Preliminaries

Let $\mathcal{T}$ be, as above, a set of n δ-fat triangles in the plane in general position. We denote the (boundary) complexity of the union of the triangles in $\mathcal{T}$ by $BC(\mathcal{T})$ and measure it by the number of edges of the union boundary. We denote the number of *holes* in that union (that is, connected components of its complement) by $H(\mathcal{T})$.

Our analysis makes use of the following two known basic technical results.

Lemma 2.1 (Combination Lemma; Edelsbrunner *et al.* [13])**.** *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two families of n_1 and n_2 triangles, respectively. Then we have*

$$BC(\mathcal{T}_1 \cup \mathcal{T}_2) \le BC(\mathcal{T}_1) + BC(\mathcal{T}_2) + O(n_1 + n_2 + H(\mathcal{T}_1 \cup \mathcal{T}_2)).$$

Remarks. (1) Spelled out, this says that the complexity of the union of the triangles in $\mathcal{T}_1 \cup \mathcal{T}_2$ is at most the sum of the complexities of the unions of the individual families plus an overhead term proportional to the number of holes in the union of the combined families and to the overall number of triangles.
(2) The lemma does not assume fatness, and thus can be applied for arbitrary triangles. (In fact, the formulation in [13] is even more general, and relates to arbitrary arrangements of segments and to any subset of cells in the overlay of two sub-arrangements.)
(3) The terms $BC(\mathcal{T}_1)$, $BC(\mathcal{T}_2)$ appear in the above bound without any multiplicative factors. This crucial property is exploited by our analysis—see below. (It fails in more general settings involving curved arcs rather than straight line segments.)

Lemma 2.2 (Matoušek *et al.* [30])**.** *For any fixed $\delta > 0$, every family of n δ-fat triangles in the plane has $O(n)$ holes, with a constant of proportionality that depends on δ.*

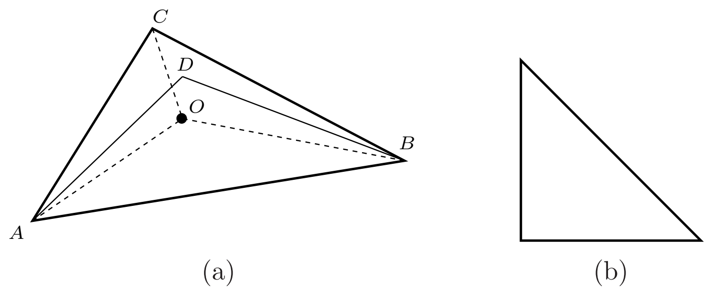

Figure 1: (a) The semi-canonization step. The triangle ABC is covered by three triangles, each of which contains the center O of the inscribed circle of ABC and has two edge orientations that are taken from a fixed set of $O(1/\delta)$ directions. Only one of these triangles (ABD) is depicted in the figure. (b) A semi-canonical right triangle, after an appropriate affine transformation.

The semi-canonization step. Following the analysis of [30], we cover each δ-fat triangle $T \in \mathcal{T}$ by a triple of "semi-canonical" $(\delta/2)$-fat triangles, each of which has a pair of edges with orientations taken from a fixed finite set $\mathcal{D}$ of $O(1/\delta)$ directions, and a third edge that also bounds T; see [30, Lemma 3.2] and Figure 1(a).

This "semi-canonization" step yields a constant number κ ($O(1/\delta^3)$, to be precise) of semi-canonical subfamilies $\mathcal{T}_1, \ldots, \mathcal{T}_\kappa$ of $(\delta/2)$-fat triangles, where the triangles in each subfamily have two edges at fixed orientations in $\mathcal{D}$, and a third edge at an orientation close to a fixed one. By construction, the union of the triangles in $\mathcal{T}$ is equal to the union of the triangles in $\mathcal{T}_1 \cup \cdots \cup \mathcal{T}_\kappa$. Putting $n_i = |\mathcal{T}_i|$, for $i = 1, \ldots, \kappa$, we have $\sum_{i=1}^{\kappa} n_i = 3n$ because each triangle is replaced by three semi-canonical triangles.

We show below that the complexity of the union of the triangles of the ith family $\mathcal{T}_i$ is $O(n_i 2^{\alpha(n_i)} \log^* n_i)$. We then note that any vertex of the union is also a vertex of the union of at most two of the sub-families (namely, those containing the triangles whose boundaries meet at the vertex). Hence the complexity of the union is upper bounded by the sum of the complexities of the unions of pairs of sub-families. Each of the latter complexities can be bounded by applying the Combination Lemma (Lemma 2.1) to the corresponding pair of families. Since the number of families is constant, we readily obtain the bound $O(n 2^{\alpha(n)} \log^* n)$ on the complexity of the union of $\mathcal{T}$.

Thus, in what follows we focus on a fixed semi-canonical family $\mathcal{T}_i$, and, to simplify the notation, we rename it $\mathcal{F}$, and denote its size, with a slight abuse of notation, by n. As in [30], by applying an appropriate affine transformation, we may assume that each triangle $T \in \mathcal{T}$ is an "almost isosceles" right triangle with one horizontal edge and one vertical edge, which meet at the lower-left vertex of T, so that the orientation of the hypotenuse of T differs from 135° by at most one degree, say; see Figure 1(b). (In particular, from now on, we ignore the dependence on δ of the constants of proportionality in our bounds.) Recall that our goal is to show that the complexity of the union of $\mathcal{F}$ is $O(n 2^{\alpha(n)} \log^* n)$, in the abused notation.

2.2 The Analysis

As shown in [30], using an interval-tree approach, we can decompose $\mathcal{F}$ into $k := O(\log n)$ subfamilies $\mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_k$, each of which corresponds to a fixed level of the tree, so that all the triangles in a fixed subfamily are stabbed by a collection of horizontal lines, each triangle in the subfamily is stabbed by precisely one line of the collection, and two triangles stabbed by different lines are disjoint (in fact, their y-projections are disjoint); see Figure 2(a). Moreover, as shown in [30], the

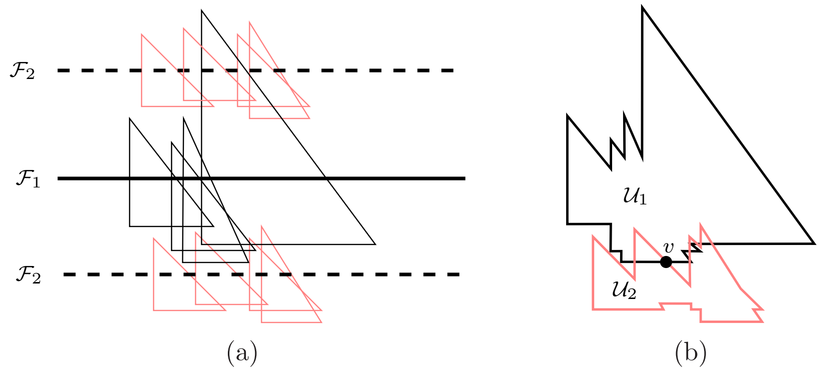

Figure 2: (a) An interval tree with two levels partitions the triangles into the subfamilies $\mathcal{F}_1$ (of the triangles stabbed by the center line) and $\mathcal{F}_2$ (of the triangles stabbed by one of the dashed lines but not by the center line). (b) Any intersection vertex v of $\partial\mathcal{U}_1$ and $\partial\mathcal{U}_2$ appears on the boundary of their union.

complexity of the union $\mathcal{U}_j$ of each subfamily $\mathcal{F}_j$ is $O(n_j 2^{\alpha(n_j)})$, where $n_j = |\mathcal{F}_j|$. To simplify the notation, write $\beta(m) = O(2^{\alpha(m)})$, with an appropriate constant of proportionality, and upper bound the complexity of $\mathcal{U}_j$ (the number of edges forming its boundary) by $n_j\beta(n_j) \le n_j\beta(n)$, $j = 1, \ldots, k$.

We next replace the triangles in each $\mathcal{F}_j$ by the collection E_j of segments that constitute the boundary of the corresponding union $\mathcal{U}_j$ (so we eliminate all portions of triangle boundaries lying in the interior of $\mathcal{U}_j$). We thus have $|E_j| \le n_j\beta(n)$. Put $E := \bigcup_{j=1}^{k} E_j$, $N := |E| \le n\beta(n)$. We may also assume that every triangle contributes at least one segment to E, so that $N \ge n$. (Triangles that fail to do so can be removed from $\mathcal{T}$, which can only increase the complexity of the union—see below for similar arguments.) Construct the arrangement $\mathcal{A}(E)$, and consider a vertex v of $\mathcal{A}(E)$, formed by the intersection of two segments of E, from two distinct sub-unions $\mathcal{U}_{j_1}, \mathcal{U}_{j_2}$ (two segments of the same sub-union can meet only at a common endpoint, and the overall number of such endpoints is $|E|$). It is easy to verify that v must be a vertex of the union $\mathcal{U}_{j_1} \cup \mathcal{U}_{j_2}$; see Figure 2(b). Using the Combination Lemma (Lemma 2.1) and the fact that this union has only $O(n_{j_1} + n_{j_2})$ holes (by Lemma 2.2), we conclude that the complexity of this union is at most $O((n_{j_1} + n_{j_2})\beta(n))$. Summing over all pairs of unions, and noting that the number of unions is only $k = O(\log n)$, we conclude that the complexity of $\mathcal{A}(E)$ is

$$X := O\left(\sum_{1\le j_1<j_2\le k} (n_{j_1} + n_{j_2})\right)\beta(n) = O\left(\sum_{1\le j_1\le k} (n_{j_1}k + n)\right)\beta(n) = O(nk\beta(n)) = O(n\beta(n)\log n).$$

The problem decomposition. We next construct a $(1/r)$-*cutting* Ξ of $\mathcal{A}(E)$, for $r = N^2/X$ (for simplicity, and with no real effect on the analysis, we ignore rounding issues), following the standard approach of [6]. We spell out the fairly standard details of the construction, to clarify that it also works for our special complexity-sensitive choice of r (a similar idea has recently been used in [20]). Specifically, we draw a random sample R of segments of E, where each segment is chosen independently with probability $p = r/N$, construct the arrangement $\mathcal{A}(R)$ and its *vertical decomposition* (see [32]). This creates a tiling of the plane into trapezoids. The expected number of trapezoids is proportional to the expected number of vertices in $\mathcal{A}(R)$, which, by our choice of r, is

$$O(pN + p^2X) = O\left(r + \left(\frac{r}{N}\right)^2 X\right) = O(r).$$

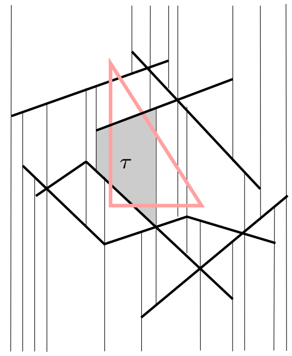

Figure 3: The vertical decomposition of an arrangement of segments. The cell τ is the only cell of the decomposition which is met by all three edges of the light-colored triangle.

Fix a cell $\tau \in \Xi$. Let E_τ denote the subset of segments in $E \setminus R$ that *cross* τ (i.e., intersect the interior of τ). It is easy to verify that on average (in a sense that we do not make too precise here) $|E_\tau| = O(N/r)$. Moreover, using the random sampling technique of Clarkson and Shor [9] (see also [23]), it can also be shown that, with high probability, $|E_\tau| = O((N/r)\log r)$ for every τ. To ensure that each cell is crossed by at most N/r segments, we take each "heavy" cell τ, namely, cells with $|E_\tau| > N/r$, and apply the resampling technique of [1, 6, 28] within τ, to refine it into subtrapezoids, so that each of them is crossed by at most N/r segments of E, and so that the overall number of trapezoids remains $O(r)$. The technical details of this step (mostly routine but adapted to our special choice of r) are spelled out in Appendix A. We continue to denote by Ξ the resulting refined cutting.

The recursive scheme. In what follows we consider intersections between triangles of $\mathcal{F}$ and trapezoids of Ξ. Such an intersection of a triangle T and a cell τ can be either *visible*, when at least one segment of E lying on the boundary of T crosses τ, or *hidden*, when there are no such segments. Hidden intersections are not handled by our analysis; in particular, we cannot bound the number of triangles intersecting a fixed trapezoid in a hidden manner by N/r. A hidden intersection between T and τ occurs, for example, when T fully contains τ. However, it can also arise in cases where some edges of T cross τ but do not appear on the boundary of the union of the subfamily $\mathcal{F}_j$ containing T, within τ. For example, it can happen that T participates in the union $\mathcal{U}_j$ of some subfamily $\mathcal{F}_j$ (clipped to τ), in the sense that its removal would change the union, but T does not appear on the union boundary. In what follows we simply ignore hidden intersections. This does not harm our analysis because within τ, each vertex v of the full union $\mathcal{U}$ must be either an endpoint of some segment of E or an intersection between two segments of E, so that these segments both cross τ. In particular, v is also a vertex of the union of the triangles that intersect τ visibly. Hence, within τ, the complexity of the union of the visible triangles is an upper bound on the complexity of the union of the entire collection.

Let $f(n)$ denote an upper bound on the complexity of the union of n fat triangles in a fixed semi-canonical family; as mentioned earlier, we may ignore the dependence of this complexity on the fatness parameter δ. Using our $(1/r)$-cutting Ξ, we derive the following recurrence for f.

For each trapezoid τ of Ξ, let n_τ denote the number of triangles T which (visibly) meet τ so that all three of their edges cross τ (and at least one of them does so in a visible manner). We refer to such triangles as *real triangles* within τ. Let m_τ denote the number of the other triangles T which (visibly) meet τ. That is, at most two of the edges of T intersect τ (and at least one of these intersections is visible); we refer to such triangles as *wedge triangles* within τ, because the intersection of any such triangle T with τ is equal to the intersection with τ of a $(\delta/2)$-fat wedge or a halfplane containing T. We note that a triangle can be a real triangle within at most one trapezoid τ. Indeed, if there were two such trapezoids we would have obtained an impossible plane embedding of $K_{3,3}$; see Figure 3 and [21] for a similar argument. In particular, $\sum_\tau n_\tau \le n$. For wedge triangles the slightly weaker bound $\sum_\tau m_\tau = O(N)$ applies, which follows since $m_\tau \le N/r$ for every τ and the number of trapezoids is $O(r)$. Similarly, $n_\tau \le N/r \le \beta(n)\log n$ for every τ, since by construction we have $N/r = X/N \le \beta(n)\log n$.

Within each trapezoid τ of our cutting Ξ, the complexity of the union can be bounded as follows. The complexity of the union of the (full) real triangles meeting τ is at most $f(n_\tau)$ (for the time being, we take the full triangles, and do not clip them to τ yet). The complexity of the union of the (full) wedges supporting the wedge triangles meeting τ is $O(m_\tau)$ [16], and the number of holes of the union of both (full) real triangles and wedges (meeting τ) is $O(m_\tau + n_\tau)$, by Lemma 2.2. (Recall that we only handle visible intersections with τ and that the ignored hidden intersections do not harm the analysis.) In order to clip the union to τ, we add to the set of the wedge triangles the four respective halfplanes bounding the four edges of τ, each of which is oppositely oriented to τ. This does not increase the asymptotic union bounds stated above. It now follows from the Combination Lemma (Lemma 2.1) that the complexity of the union within τ is at most $f(n_\tau)+O(m_\tau+n_\tau)$. (Recall the important property that $f(n_\tau)$ appears in the bound without any additional multiplicative factor; see Remark (3) following Lemma 2.1.) Hence, summing these bounds over all trapezoids, we obtain:

$$f(n) \le \sum_\tau \big(f(n_\tau) + O(m_\tau + n_\tau)\big) = \sum_\tau f(n_\tau) + O(n\beta(n)). \tag{1}$$

We next unroll the recurrence and stop as soon as $n_\tau \le c$, where $c > 0$ is some absolute constant. Note that, at any fixed level of the recursion, the sum $\sum_\tau n_\tau$, over all trapezoids of all subproblems, is at most (the initial value of) n, so the corresponding sum of the overhead terms, at any fixed level, is $O(n\beta(n))$. Since the maximum problem size n_h at each level $h \ge 2$ is only $O(\beta(n_{h-1})\log n_{h-1})$, the number of levels in the recurrence is $O(\log^* n)$, which yields

$$f(n) = O(n\beta(n)\log^* n) = O(n2^{\alpha(n)}\log^* n).$$

That is, we have obtained the following main result:

Theorem 2.3. *The complexity of the union of n δ-fat triangles in the plane is $O(n2^{\alpha(n)}\log^* n)$, with a constant of proportionality depending on δ.*

Similar results can be obtained for any collection of polygons in the plane that can be decomposed into, or covered by, a total of n δ-fat triangles. In particular, Van Kreveld [33] has shown that this property holds for collections of polygons with n vertices in total, none of which has a δ-*corridor*, that is, a passage between two edges, such that the ratio between its width and length is less than δ. Thus we obtain:

Theorem 2.4. *The complexity of the union of polygons in the plane, none of which has a δ-corridor, with a total of n vertices, is $O(n2^{\alpha(n)}\log^* n)$, with a constant of proportionality depending on δ.*

Remarks. (1) We note that a key step of the analysis is reducing the number of intersections among the triangle boundaries that may eventually appear as vertices of the union from $O(n^2)$ (which is just the arrangement complexity) to $X = O(n\beta(n)\log n)$, with only a slight increase in the number of segments (that is, from $O(n)$ to $N \le n\beta(n)$). This has eventually led to the decomposition of the problem into roughly $O(N/\log n)$ subproblems, each of size close to $O(\log n)$, which facilitated the derivation and efficient solution of recurrence (1).

(2) Investigating the dependence of the complexity of the union on δ is a refinement of the current analysis which we plan to carry out in the full version of the paper.

3 Applications

Set cover and ε-nets. A *range space* $(P, \mathcal{R})$ is a pair consisting of an underlying universe P of objects, and a certain collection $\mathcal{R} \subseteq 2^P$ of its subsets (*ranges*). We assume in this discussion that P and $\mathcal{R}$ are finite. The *set cover* problem for $(P, \mathcal{R})$ is to find a minimum-size subcollection $S \subseteq \mathcal{R}$ whose union covers P. Put $\text{OPT} := |S|$. The general problem is known to be NP-hard and the problem remains NP-hard in most geometric settings; see the discussion in [3] and the references therein. The standard greedy algorithm [7] yields a set cover with an approximation factor $O(\log |P|)$ without any further assumptions on the range space. This factor is (asymptotically) best possible to compute in polynomial time, under appropriate complexity-theoretic assumptions [4, 22]. However, as observed in [5, 18] (see also [8]), the approximation factor can be improved to $O(\log \text{OPT})$, still achievable in expected polynomial time, for many geometric scenarios.

This improvement is closely related to the notion of *ε-nets* in a "dual" variant of such a range space. Specifically, in this dual context, a subset $N \subseteq \mathcal{R}$ is called a *dual ε-net* for $(P, \mathcal{R})$, for some given $0 < \varepsilon < 1$, if every point $p \in P$ which is contained in more than $\varepsilon|\mathcal{R}|$ ranges of $\mathcal{R}$ is contained in a range of N. The classical result of Haussler and Welzl [23], specialized to this context, asserts that, if $(P, \mathcal{R})$ has so-called *finite VC-dimension*, then it admits dual ε-nets of the above kind of size $O((1/\varepsilon) \log(1/\varepsilon))$, where the constant of proportionality depends on the VC-dimension.

As shown in [5, 18], if $(P, \mathcal{R})$ admits dual ε-nets of size $O((1/\varepsilon)\varphi(1/\varepsilon))$, then one can obtain approximation factor $O(\varphi(\text{OPT}))$ in expected polynomial time for the corresponding set cover problem. Thus, since we always have $\varphi(1/\varepsilon) = O(\log(1/\varepsilon))$, the "default" result, for range spaces of finite VC-dimension, is an approximation factor $O(\log \text{OPT})$, as mentioned above.

As observed in several recent works [3, 10, 34], the size of dual ε-nets is strongly related to the complexity of the union of the given ranges. Improving upon a slightly weaker bound given in [10], the following result was established by the authors in [3] (see also [34]). Let P be a finite set of points in the plane, and let $\mathcal{R}$ be a collection of planar regions of *constant description complexity* (meaning that they are semi-algebraic sets defined in terms of a constant number of polynomial equations and inequalities of constant maximum degree), such that the complexity of the union of any subset of r of these regions is $O(r\phi(r))$. Then there exist dual ε-nets for $(P, \mathcal{R})$ whose size is $O((1/\varepsilon) \log \phi(1/\varepsilon))$. This in turn yields the approximation factor (still computable in expected polynomial time) $O(\log \phi(\text{OPT}))$ for the corresponding set cover problem. Incorporating the bound of Theorem 2.3 we obtain the following two corollaries.

Theorem 3.1. *Any dual range space of δ-fat triangles and points in the plane admits an ε-net of $O((1/\varepsilon) \log \log^* (1/\varepsilon))$ triangles, for any $0 < \varepsilon \leq 1$.*

Theorem 3.2. *There exists a randomized expected polynomial-time algorithm that, given a set P of points in the plane and a set $\mathcal{T}$ of δ-fat triangles that cover P, computes a set cover $\mathcal{T}' \subseteq \mathcal{T}$ for P of size $O(\text{OPT} \log \log^* \text{OPT})$, where OPT is the size of the smallest such cover.*

Hidden surface removal. In a typical hidden surface removal problem, we are given a set of n pairwise-disjoint objects in 3-space, and a viewing point v, and the goal is to construct the view of the given scene, as seen from v. This view consists of a subdivision of the viewing plane into maximal connected regions in each of which at most one object (or some portion thereof) can be seen. The resulting structure is the so-called *visibility map* of the given objects as seen from v.

Katz *et al.* [24] have proposed an output-sensitive algorithm to construct the visibility map, whose running time is $O((U(n) + k)\log^2 n)$, where $U(r)$ is an upper bound on the complexity of the union of the projections to the viewing plane of any r of the input objects (here $U(\cdot)$ is also assumed to be a super-additive function), and k is the complexity of the visibility map. The space requirement in this case is $O(U(n)\log n)$.

When the objects are horizontal δ-fat triangles, and the viewing plane is also horizontal, the projection of the input triangles are also δ-fat. Hence, applying the bound of Theorem 2.3, we obtain the following improved bounds on the performance of the algorithm in [24].

Theorem 3.3. *Given a set of n pairwise-disjoint horizontal δ-fat triangles in 3-space, the view of this set from any fixed point v can be computed in $O((n2^{\alpha(n)}\log^* n + k)\log^2 n)$ time, using $O(n2^{\alpha(n)}\log^* n \log n)$ space, where k is the complexity of the visibility map.*

4 Concluding Remarks

Our bound on the complexity of the union of δ-fat triangles is the first improvement over the old bound of [29, 30]. There is still a (smaller) gap between this bound and the known lower bound, and closing this gap remains a hard challenge. It is interesting to note that the new bound moves one step up the inverse-Ackermann hierarchy. We hope that the new ideas presented in this paper can be enhanced to further improve the bound up the hierarchy, ultimately reaching a bound of the form $O(nG(\alpha(n)))$, for some elementary function G. This is what we conjecture at the moment. Improving the upper bound all the way down to the known lower bound $\Omega(n\alpha(n))$ seems a much harder challenge.

Acknowledgments. We wish to thank Sariel Har-Peled for numerous helpful discussions and Mark de Berg for several useful suggestions.

References

[1] P. K. Agarwal, J. Matoušek, and O. Schwarzkopf, Computing many faces in arrangements of lines and segments, *SIAM J. Comput.*, 27:491–505, 1998.

[2] P.K. Agarwal, J. Pach, and M. Sharir, State of the union (of geometric objects), In *Proc. Joint Summer Research Conf. on Discrete and Computational Geometry: 20 Years Later*, Contemp. Math. 452, Amer. Math. Society, pp. 9–48, 2008.

[3] B. Aronov, E. Ezra, and M. Sharir, Small-size ε-nets for axis-parallel rectangles and boxes, *SIAM J. Comput.*, in press. An extended abtract appeared in *Proc. 41th Annu. ACM Symp. Theory Comput.*, pp. 639–648, 2009.

[4] M. Bellare, S. Goldwasser, G. Lund, and A. Russell, Efficient probabilistically checkable proofs and applications to approximation, In *Proc. 25th Annu. ACM Symp. Theory Comput.*, pp. 294–304, 1993.

[5] H. Brönnimann and M. T. Goodrich, Almost optimal set covers in finite VC dimensions, *Discrete Comput. Geom.*, 14:463–479, 1995.

[6] B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica*, 10:229–249, 1990.

[7] V. Chvátal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.

[8] K. L. Clarkson, Algorithms for polytope covering and approximation, In *Proc. 3rd Workshop on Algorithms and Data Structures*, pp. 246–252, 1993.

[9] K. L. Clarkson and P. W. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.*, 4:387–421, 1989.

[10] K. L. Clarkson and K. Varadarajan, Improved approximation algorithms for geometric set cover, *Discrete Comput. Geom.*, 37:43–58, 2007.

[11] M. de Berg, Improved bounds on the union complexity of fat objects, *Discrete Comput. Geom.*, 40(1):127–140, 2008.

[12] M. de Berg, Better bounds on the union complexity of locally fat objects, In *Proc. 26th Annu. Sympos. Comput. Geom.*, pp. 39–47, 2010.

[13] H. Edelsbrunner, L. J. Guibas, and M. Sharir, The complexity and construction of many faces in arrangements of lines and of segments, *Discrete Comput. Geom.*, 5(2):161–196, 1990.

[14] A. Efrat, The complexity of the union of (α, β)-covered objects, *SIAM J. Comput.*, 34:775–787, 2005.

[15] A. Efrat and M. Katz, On the union of κ-curved objects, *Comput. Geom. Theory Appl.* 14:241–254, 1999.

[16] A. Efrat, G. Rote, and M. Sharir, On the union of fat wedges and separating a collection of segments by a line, *Comput. Geom. Theory Appls.* 3:277–288, 1993.

[17] A. Efrat and M. Sharir, The complexity of the union of fat objects in the plane, *Discrete Comput. Geom.* 23:171–189, 2000.

[18] G. Even, D. Rawitz, and S. Shahar, Hitting sets when the VC-dimension is small, *Inform. Process. Letts.*, 95:358–362, 2005.

[19] E. Ezra, On the union of cylinders in three dimensions, In *Proc. 49nd Annu. IEEE Sympos. Found. Comput. Sci.*, pp. 179–188, 2008.

[20] E. Ezra and J. Matoušek, Spanning trees with low crossing numbers for line-segments, *Manuscript,* 2010.

[21] E. Ezra and M. Sharir, Almost tight bound for the union of fat tetrahedra in three dimensions, *J. ACM* 57(1), 2009.

[22] U. Feige, A threshold of $\ln n$ for approximating set cover, *J. ACM*, 45(4):634–652, 1998.

[23] D. Haussler and E. Welzl, ε-nets and simplex range queries, *Discrete Comput. Geom.*, 2:127–151, 1987.

[24] M. J. Katz, M. H. Overmars, and M. Sharir, Efficient hidden surface removal for objects with small union size, *Comput. Geom. Theory Appls.*, 2(4):223–234, 1992.

[25] 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.*, 1:59–71, 1986.

[26] J. Matoušek, *Lectures on Discrete Geometry*, Springer-Verlag New York, 2002.

[27] J. Matoušek, Efficient partition trees, *Discrete Comput. Geom.*, 8:315–334, 1992.

[28] J. Matoušek, Reporting points in halfspaces, *Comput. Geom. Theory Appl.* 2:169–186, 1992.

[29] J. Matoušek, N. Miller, J. Pach, M. Sharir, S. Sifrony, and E. Welzl, Fat triangles determine linearly many holes, In *Proc. 32nd Annu. IEEE Sympos. Found. Comput. Sci.*, pp. 49–58, 1991.

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

[31] J. Pach and G. Tardos, On the boundary complexity of the union of fat triangles, *SIAM J. Comput.*, 31:1745–1760, 2002.

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

[33] M. van Kreveld, On fat partitioning, fat covering and the union size of polygons, *Comput. Geom. Theory Appl.*, 9(4):197–210, 1998.

[34] K. Varadarajan, Epsilon nets and union complexity, *Proc. 25th Annu. Sympos. Comput. Geom.*, pp. 11–16, 2009.

A Appendix - The Refined Cutting

The refined decomposition is constructed in the following standard manner. Define the *weight factor* t_τ of a trapezoid τ, generated at the first stage of the construction, to be $\lceil |E_\tau| \cdot r/N \rceil$. Trapezoids τ with $t_\tau \le 1$ remain intact, and are included in the refined cutting. If $t = t_\tau > 1$, we pick a $(1/t)$-net N_τ for E_τ of size $ct \log t$, where $c > 0$ is a sufficiently large absolute constant (whose existence is guaranteed by the standard ε-net theory [23]). We then construct the arragement of the segments in N_τ (clipped to within τ), decompose each of the cells in the resulting arrangement within τ into trapezoidal subcells, and add them all to the refined cutting. The overall number of new cells within τ is $O(t^2 \log^2 t)$. Since N_τ is a $(1/t)$-net (see [23]), it follows that each new subcell meets at most $|E_\tau|/t \le N/r$ segments of E_τ (and thus of E). It remains to show that the expected number of new subcells remains $O(r)$.

To this end, for a subset $Q \subseteq E$, let $\mathrm{VD}(Q)$ be the set of trapezoidal cells in the vertical decompostion of the arrangement $\mathcal{A}(Q)$, and let $\mathrm{VD}_t(Q)$ be the subset of these cells τ with $t_\tau \ge t$, for a positive parameter t. Following the analysis in [1, 6, 28], let $\mathcal{D} = \bigcup_{Q \subseteq E} \mathrm{VD}(Q)$ be the set of all trapezoidal cells defined by all possible subsets of E. The *defining set* $D(\tau)$ of a cell τ in $\mathcal{D}$ consists of the at most four segments in E that define τ, and its *killing set* $K(\tau)$ is the set of segments in E that cross τ (see [1, 6, 28] for a more detailed description). As in previous studies, the following "axiom" is easily seen to hold: A cell $\tau \in \mathcal{D}$ belongs to $\mathrm{VD}(R)$ if and only if $D(\tau) \subseteq R$ and $K(\tau) \cap R = \emptyset$. Using the analysis in [1, 6, 28] we obtain the following lemma (where $p = r/N$ is the probability of choosing a segment in R):

Lemma A.1 (Exponential Decay Lemma; Agarwal *et al.* [1])**.** *Let R, E, and t be defined as above, then*

$$\mathbf{Exp}\left\{|\,VD_t(R)|\right\} = O\left(2^{-t}\,\mathbf{Exp}\left\{|\,VD(R')|\right\}\right),$$

where $R' \subseteq E$ is a random sample of E, where each segment is chosen into R' independently with probablity $p' := p/t = r/(Nt)$.

Applying the bound in Lemma A.1, we have:

$$\textbf{Exp}\left\{|\mathrm{VD}_t(R)|\right\} = O\left(\frac{p^2X + tpN}{t^2 2^t}\right). \tag{2}$$

We next bound the expected number of cells in the refined cutting using (2). As noted above, the expected number of intact cells is $O(r)$. We bound below the expected number of cells generated at the second sampling step (where $c' > 0$ below is a sufficiently large constant):

$$\textbf{Exp}\left\{\sum_{t\geq 1} \sum_{\substack{\tau\in\mathrm{VD}(R)\\ t\leq t_\tau\leq t+1}} c'\cdot t_\tau^2 \log^2 t_\tau\right\} = \sum_{t\geq 1}\textbf{Exp}\left\{\sum_{\substack{\tau\in\mathrm{VD}(R)\\ t\leq t_\tau\leq t+1}} c'\cdot t_\tau^2 \log^2 t_\tau\right\} =$$

$$O\left(\sum_{t\geq 1} t^2\log^2 t\cdot\frac{p^2X + tpN}{t^2 2^t}\right) = O(p^2X + pN) = O(r).$$